

Is Attention Both Necessary and Sufficient for Consciousness?

Antonios Kaldas

MB.,BS.(Syd), GradDipPhil(MQ), MRes(MQ).

*A thesis submitted in fulfillment
of the requirements for the degree
of Doctor of Philosophy*

Department of Philosophy
Macquarie University

1st March 2019

Principal Supervisors:
Dr Colin Klein (2016-2017)
Prof Richard Menary (2018-2019)

Adjunct Supervisor:
A/Prof Alex Holcombe (University of Sydney)

Table of Contents

LIST OF TABLES	5
LIST OF FIGURES.....	7
ABSTRACT	9
ACKNOWLEDGMENTS.....	12
1 Q: IS ATTENTION BOTH NECESSARY AND SUFFICIENT FOR CONSCIOUSNESS?	15
1.1 INTRODUCTION	15
1.2 WHY DOES Q MATTER?	18
1.3 ARGUMENTS—TWO PAPERS.....	20
1.4 ANALYSIS	23
1.5 OUTLINE OF THESIS	38
 <i>PART I. (RE-)LAYING THE FOUNDATIONS</i>	
2 PHENOMENAL CONSCIOUSNESS.....	45
2.1 PHENOMENAL CONSCIOUSNESS AND Q.....	45
2.2 A BRIEF HISTORY OF “CONSCIOUSNESS”	46
2.3 THE ONTOLOGY OF CONSCIOUSNESS.....	48
2.4 DEFINING PHENOMENAL CONSCIOUSNESS.....	50
2.5 THE <i>CONTENT-PHENOMENALITY</i> DISTINCTION.....	57
2.6 CHAPTER SUMMARY.....	59
3 THE MANY FACES OF ATTENTION.....	61
3.1 A BRIEF HISTORY OF ATTENTION.....	61
3.2 PRELIMINARIES TO DEFINING ATTENTION	63
3.3 A TAXONOMY OF ATTENTION.....	75
3.4 MY APPROACH TO DEFINING ATTENTION.....	103
3.5 CHAPTER SUMMARY.....	108
4 RELATIONSHIPS	109
4.1 INTRODUCTION	109
4.2 THE SET THEORETICAL FRAMEWORK (STF).....	113
4.3 WHAT OTHERS HAVE SAID.....	134
4.4 KINDS OF RELATIONSHIP.....	150
4.5 CHAPTER SUMMARY.....	155
5 WORKING MEMORY	157
5.1 INTRODUCING WORKING MEMORY (WM)	157
5.2 HISTORY	158
5.3 WHAT IS WORKING MEMORY?	161
5.4 THE DUAL-ASPECT NATURE OF WORKING MEMORY	176
5.5 BOUNDARIES OF WORKING MEMORY?	180
5.6 CHAPTER SUMMARY.....	181

PART II. ADDRESSING Q

6	PULLING ATTENTION AND CONSCIOUSNESS APART	185
6.1	RECAP AND OUTLINES	185
6.2	ATTENTION WITHOUT CONSCIOUSNESS.....	188
6.3	CONSCIOUSNESS WITHOUT ATTENTION.....	195
6.4	CHAPTER SUMMARY.....	225
7	PHENOMENAL OVERFLOW	227
7.1	THE ENIGMA OF PHENOMENAL OVERFLOW	227
7.2	SOME MORE FOUNDATIONS	229
7.3	FLORENCE AND BEN.....	241
7.4	EPISTEMIC WOES	245
7.5	CHAPTER SUMMARY.....	266
8	TRIANGULATING CAPACITY LIMITATIONS.....	269
8.1	COMPARING CAPACITIES	269
8.2	THE CAPACITY OF WORKING MEMORY	270
8.3	THE CAPACITY OF ATTENTION.....	279
8.4	THE CAPACITY OF CONSCIOUSNESS.....	285
8.5	THE WITCHES' HAT MODEL.....	293
8.6	ALTERNATIVE INTERPRETATIONS	299
8.7	CHAPTER SUMMARY.....	312
9	ANSWER(S) TO Q	313
9.1	CONCLUSIONS, IMPLICATIONS, APPLICATIONS	313
9.2	SEEKING ANSWERS	314
9.3	WORKING MEMORY	325
9.4	IMPLICATIONS FOR THE CURRENT DISCOURSE ON ATTENTION AND CONSCIOUSNESS	327
9.5	BROADER IMPLICATIONS AND APPLICATIONS	332
9.6	FINAL REFLECTIONS	353
	APPENDICES	355
	APPENDIX 1. DISTINCTIONS BETWEEN DIFFERENT KINDS OF ATTENTION.....	355
	APPENDIX 2. FOUR UNLIKELY SCENARIOS	360
	APPENDIX 3. TABULATED COMPARISONS OF STF WITH OTHER FRAMEWORKS	364
	APPENDIX 4. CLASSIFICATION OF AUTHORS BY STF SCENARIO	372
	APPENDIX 5. MODELS OF WM: A SELECTION	385
	APPENDIX 6. METAPHORS OF WM: A SELECTION	391
	APPENDIX 7. ANATOMY AND PHYSIOLOGY OF FOVEAL VISION	394
	BIBLIOGRAPHY	399

List of Tables

Table 1. A summary of some important distinctions among kinds of attention. (3.2.3)

Table 2. A definitional taxonomy of attention. (3.3.1)

Table 3. The Matrix of all possible STF Scenarios. (4.2.3.1)

Table 4. Four kinds of relationship between attention and consciousness. (4.4.1)

Table 5. Four unlikely STF Scenarios. (Appendix 2)

Table 6. Comparison of STF with the descriptions of Iwasaki (1993, p. 212).
(Appendix 3)

Table 7. Comparison of STF with the models of Lamme. (Appendix 3)

Table 8. Comparison of STF with the categories of Schwitzgebel (2007). (Appendix
3)

Table 9. Comparison of STF with the framework of Koch and Tsuchiya (2007).
(Appendix 3)

Table 10. Comparison of STF with the propositions of De Brigard (2010, pp. 189–
190). (Appendix 3)

Table 11. Comparison of STF with the CAD framework of Montemayor & Haladjian
(2015, pp. 5–6). (Appendix 3)

Table 12. Comparison of STF with the descriptions of Pitts et al., (2018, p. 2).
(Appendix 3)

Table 13. Classification of Authors by Scenario. (Appendix 4)

Table 14. Alphabetical list of Scenarios by author. (Appendix 4)

Table 15. Models of Working Memory. (Appendix 5)

Table 16. Metaphors of Working Memory. (Appendix 6)

Table 17. Calculations of foveal and foveolar visual fields. (Appendix 7)

List of Figures

Figure 1. The four possible Combinations of attention and consciousness. (4.2.2)

Figure 2 A graphical representation of one way that gradations in the overlap between attention and consciousness might be represented. (4.2.4.2)

Figure 3 The preponderance of philosophers vs non-philosophers among the Scenarios of STF. (4.3.6)

Figure 4. Two models of working memory. (5.3.1)

Figure 5 Witches' Hat Model of conscious content. (8.5)

Figure 6 Four models of visual awareness and its relation to attention, modified from Lamme (2003) p. 13. (Appendix 3)

Figure 7 Consciousness vs Attention, modified from Van Boxtel et al., (2010, p. 6, Figure 2a). (Appendix 3)

Figure 8 Visual field angle and diameter. (Appendix 7)

Abstract

Is attention both necessary and sufficient for consciousness? Call this central question of this treatise, “Q.” We commonly have the experience of consciously paying attention to something, but is it possible to be conscious of something you are not attending to, or to attend to something of which you are not conscious? Where might we find examples of these?

This treatise is a quest to find an answer to Q in two parts. Part I reviews the foundations upon which the discourse on Q is built. Different inputs to Q produce different answers. After consideration of the many ways “attention” and “consciousness” have been defined, I settle upon *phenomenal consciousness* and *Executive Attention* (defined as a suite of strategies for structuring cognition for further processing implemented by the executive of working memory) as the most interesting inputs to Q, and the ones on which Part II focuses.

Attention without consciousness seems relatively easy to establish empirically, but consciousness without attention is much harder. The putative candidates all seem to have major problems, but I build a strong abductive case for the hitherto ignored case of foveal phenomenal overflow. We consciously see far more detail in our foveal fields than we can Executively Attend, although there is a serious obstacle to our ever confirming that empirically—identifying conscious content relies on Executive Attentional report. Triangulating the capacity limitations of attention, consciousness, and working memory strengthens this case for consciousness without attention, and suggests that cognition must work something like my “Witches’ Hat Model,” on which content can become conscious outside of Executive Attention or working memory. I conclude with some reflections on the implications of my arguments for the discourse on Q, and for other discourses such as the ontologies of attention and consciousness, theories of consciousness, some other cognitive concepts, and ethical considerations in humans, animals, and machines.

A conclusive answer to Q continues to elude us. It may perhaps be an ultimately insoluble conundrum. But it is the very essence of humanity to seek an answer, and in so doing, to improve our understanding of our own nature:

“The proper study of mankind is man.”¹

¹ Alexander Pope, *An Essay on Man*, 2.1.

Statement of Originality

*This work has not previously been submitted for a degree or diploma in any university.
To the best of my knowledge and belief, the thesis contains no material previously
published or written by another person except where due reference is made in the
thesis itself.*

Antonios Kaldas

1st March 2019

Acknowledgments

Conducting this kind of prolonged and concentrated philosophical research is privilege, and the support of others indispensable to enjoying that privilege. My beautiful wife Dalia and wonderful children Samuel and Bethany not only granted me the time to work on it, but unfailingly encouraged me and believed that I could complete it. They provided an honest and creative sounding board when there was no one else to talk to. For their unconditional love, I shall always be grateful.

My journey in academic philosophy would have met with a premature end if not for the kind encouragement and common-sense of the late Peter Menzies, the very present Jeanette Kennett, and my Masters supervisor, Rachael Brown, from whom I learnt so much.

At the centre of a project like this, of course, the Principal Supervisor carries most of the load. I feel both honoured and blessed to have started this PhD journey with Colin Klein (now at ANU) and completed it with Richard Menary. They have both been incredibly easy to work with, patient, understanding, wise, and a fountain of knowledge, as well as being jolly decent human beings. I can never say thank you enough for the many hours you invested in my work, and the sensitive way you guided me out of my sillier notions and towards true understanding. You have taught me the art of philosophy by being models of its noblest embodiment. My Adjunct Supervisor from the world of empirical psychology, Alex Holcombe, has

been more than generous with his time and wisdom, and an invaluable asset in what were, for me, unknown lands. For all that my supervisors gave me, I shall always be grateful.

The ideas in this thesis benefitted from a number of discussions with others, including a Macquarie University Higher Degree Research Masterclass with Logi Gunnarson (March 2017), a presentation at the IACAP Annual Conference (Warsaw, June 2018), and a poster at the ASSC 22nd Annual Conference (Krakow, June 2018). The latter two were funded by a PGRF grant. My overall research project was supported with an RTF scholarship.

I also had the pleasure and privilege of refining various aspects of my thesis through feedback and discussions, briefer or longer, electronically or in person, from a number of people including (in alphabetical order) Peter Adamson, Michael Archer, David Chalmers, Martin Eimer, Ran Hassin, Patricia Hutchings, Greame Jackson, Dimitri Kokkinos, Chris Letheby, John Mitry, Lisa Miracchi, Dinyar Mistry, Anina Rich, Wendy Rogers, Jennifer Windt, and Nao Tsuchiya. Of special help were the regular meetings that came to be called “Colin’s Crowd,” in which postgraduate students under the supervision of Colin Klein took turns to present their work-in-progress and receive feedback from the others, under Colin’s gentle direction. Thank you to (in alphabetical order) Peter Clutton, Stephen Gadsby, Meg Ivory, Anneli Janssen, Dorian Minors, Trisha Nowland, and Matthew Tuxford for all I learned from your creative ideas and, especially, the issues on which we disagreed.

A special thanks to our dogs Ruby, the late Merlin, and the recent Arthur, and cat Bendy for their daily company at my feet, in my lap, and ornamenting the window sill, keeping a very close philosophical eye (sometimes two) on my work.

1 Q: Is Attention Both Necessary and Sufficient for Consciousness?

1.1 Introduction

“The relationship between attention and awareness is one of the most hotly debated issues in neuroscience and psychology”
(Cohen, Alvarez, & Nakayama, 2011, p. 1170).

You are walking along an isolated dirt track in the lush semi-tropical rainforest of an Australian National Park. Your senses are bombarded by a glorious riot of perceptions—the piercing blue-green of myriads of gum leaves strewn across a brilliant bright blue sky, punctuated by smooth grey branches. Bright green and red king parrots flit by, while the cacophonous cries of corellas are occasionally punctuated with the pure single notes of bellbirds. The damp red soil smells of life,

life gently bubbling and bounding and burgeoning, as the gentle warm breeze of spring caresses your cheeks and welcomes you into nature's warm embrace. Memories of childhood picnics well up in your heart, mixed with emotions too fast to savour—the sense of freedom, the profound privilege of being one with nature, the joy of being alive.

All this takes a long time to describe but is the experience of a single moment. It seems so simple and so natural, but it is the philosopher's task to pull it apart in order to better understand it. Gandalf would be disappointed,² but we shall adopt a more optimistic attitude and hope that dissecting experiences such as the one described will only serve to deepen them with additional wonder and joy at the marvel that is human cognition.

As you walk through the forest, you seem to *consciously experience* so many things at once. Yet you cannot *attend* to them all at once. Attend to the cry of a bird, and you cease attending to the brightness of the sky. Attend to the brightness of the sky, and you cease to attend to the softness of the breeze. Even if, by a heroic effort of will, you manage to attend to two of these things at once, you will likely find it almost impossible to attend to three. To attend to the same high degree to *everything* your senses and inner life present at once is quite certainly impossible. Yet it seems somehow wrong to say that by ceasing to attend to a thing, you thereby cease to consciously experience it. Most people, I think, would rather say something like this: "I am less aware³ of some things, and more aware of others. I may not be able to count the leaves on a branch if I am currently attending to the cry of the corellas, but I am certainly *aware* that there are leaves just there, and that they are leaves rather than ice blocks." And yet, are you really aware of those leaves? Without paying at least some attention to them, how much can you really tell me about them apart from the fact of their existence? Is it possible to be aware of a thing without being

² "He that breaks a thing to find out what it is has left the path of wisdom", JRR Tolkien, *Lord of the Rings*, Fellowship of the Ring.

³ "Awareness" and "consciousness" can mean different things, but I use "aware" in this section in a less precise folk psychological sense. I define "consciousness" much more rigorously in Chapter 2.

able to tell anything about it? Or is it the case that unless you attend to a thing, you are not in any meaningful sense truly consciously aware of that thing?

1.1.1 The Question, Q

This treatise is a quest to answer one question:

Q: Is attention both necessary and sufficient for consciousness?

This is *not* a simple question. The relationship between attention and conscious experience (consciousness for short) lies at the heart of a vigorous contemporary debate in the philosophical and cognitive scientific literature. The number and variety of views on offer reflects the complexity of Q. What do “attention” and “consciousness” mean? What constitutes evidence for necessity or sufficiency? How is the empirical evidence best interpreted? What underlies the relationship between attention and consciousness? And what does all this mean for our understanding of the mind more generally? These are the kinds of questions I address in the pages to come.

At its simplest, Q is a question about a *pattern of occurrence*. If attention is necessary for consciousness, then there can be no instance of consciousness that is not accompanied by attention. Whenever you have consciousness, necessarily, you also have attention. If attention is sufficient for consciousness, then there can be no instance of attention without consciousness. Whenever you have attention, necessarily, you also have consciousness. In other words, Q asks whether attention and consciousness can “come apart,” or whether they must always come together.

The scope of Q may be general or specific. We may ask the question generally of all cognition, whether human, animal, alien, computer, or any other kind of cognition. Exploring this general question empirically is, of course, far more difficult than

exploring the specific question as it applies to human cognition. It is, therefore, the more specific question with which I here mostly concern myself. Human cognition here signifies both normal human cognition (if there is such a thing) and abnormal human cognition—cases where cognition functions in “pathological” ways.⁴

1.1.2 Outline of This Chapter

The chief purpose of this treatise is to make some progress towards an answer to Q by building on the existing literature in new and hopefully interesting and fruitful ways. To that end, this first chapter begins with consideration of why Q matters—what is at stake. I argue that Q offers a valuable window not only into the nature of both attention and consciousness themselves, but thereby into human cognition generally. I then briefly critique two interesting papers that take opposing views on the answer to Q to lay out some principles that will guide my quest. I conclude this opening chapter with an outline of the remainder of the thesis.

1.2 Why Does Q Matter?

Why should we pursue an answer to Q? I should confess from the outset that my chief motivation in this investigation is sheer unbridled and unapologetic curiosity. Human beings are fascinating in many ways, but none so beguiling as their ability to do things like attend to stuff and have conscious experiences. In this section I raise

⁴ See 4.2.3.2 for more detail.

some very practical applications to which my research might be of use, but I am ultimately interested in *understanding*.

Understanding consciousness has been called “The Most Interesting Problem in the Universe” (Turausky, 2014). The bewildering variety of largely mutually inconsistent theories of consciousness still being discussed perhaps indicates how far we are from a solution. And yet, consciousness is what makes us who we are. We define our selfhood largely as our stream of conscious experience, as demonstrated by Parfit’s (1971) famous thought experiments. And yet, as close as our stream of conscious experience is, so far is it inscrutable to us. We can say far more—and with far greater confidence—about things going on outside our selves than we can about our selves, in themselves.

This inscrutability has inspired diverse approaches that attempt to pierce through it. From the very simple, such as the assiduous introspections of Descartes, to the very complex, employing brain scanners, to hybrid experimental paradigms that manipulate aspects of conscious experience, dancing on the gently lapping shore between the conscious and the unconscious. Consciousness is intimately interwoven with other cognitive concepts, so an eminently amenable approach is to probe consciousness through its connections. Many such connections have been and continue to be employed—connections with memory, imagination, perception, and action, to name a few. But the connection of consciousness to attention seems to be particularly promising. It is so close that it has led many to virtually identify conscious with attention (3.3.3, 4.3), in a way that has rarely occurred for any other cognitive concept.

The connection of consciousness to attention also opens up a world of empirical research that can be used to probe the nature of consciousness. Scientists (read psychologists, neuroscientists, etc., as opposed to philosophers) largely ignored consciousness for most of the twentieth century, but they studied attention assiduously. Bridging the philosophical world of consciousness studies and the scientific world of attention studies has been the interdisciplinary principle of recent decades, and promises to deliver the best of both worlds. But we must also be wary.

We cannot, for example, simply assume that consciousness and attention always go together as some have, and thereby make attention a proxy for consciousness in research. In this treatise, I provide some compelling reasons (summarised in 9.5.2.3) to avoid this trap.

Attention and consciousness do not exist in a vacuum, but are part of larger *cognitive economy*— the sum of all cognitive processes occurring in an individual subject, whether synchronically or diachronically. Exploring Q leads to all sorts of interesting insights into the things going on around attention and consciousness, so to speak, some of which I explore briefly in the last chapter, where I also highlight some other practical implications that arise from consideration of Q.

1.3 Arguments—Two Papers

I begin this quest to explore Q by analysing two recent papers on the topic. I do not claim they are foundational or central texts on the topic—only that they serve beautifully to introduce and illustrate many of the chief issues that my approach to Q will address in this treatise. I pursue a much fuller literature review encompassing a broad array of authors and positions in Chapters 3 and 4. The two papers in this section are Koch and Tsuchiya (2007) and De Brigard and Prinz (2010).

1.3.1 Koch and Tsuchiya (KT)

Representing a “no” answer to Q—arguing that attention and consciousness can and do indeed come apart, is Koch and Tsuchiya (2007) (henceforth referred to as KT). It

is an empirically rich paper, with some careful, although not always conclusive arguments. I break down the overall argument of KT in four propositions, in the order in which KT themselves present them:

KT1. *Attention and consciousness have “substantially different functions.”*

KT2. *There are empirical instances where a subject can be voluntarily turning her attention to an object⁵ without simultaneously being consciously aware of that object.*

KT3. *There are empirical instances where a subject can be consciously aware of an object without simultaneously voluntarily turning her attention to that object (or at least, hardly turning at all—the “near-absence of attention”).*

KT4. *Attention and consciousness can sometimes have opposing effects.*

These four propositions together constitute, KT suggest, a powerful argument that attention is neither necessary nor sufficient for consciousness. It is worth asking what role each of these propositions is playing. On careful consideration, KT1 and KT4 do not unequivocally negate Q, but provide evidence that is more difficult—but not impossible—to reconcile with a positive answer. On KT1, a single process may subserve different functions, and these functions may even have opposing effects (KT4), given the right circumstances. On the other hand, KT2 and KT3, if confirmed, would seem to provide an unequivocal answer of no to Q (although the qualification at the end of KT3 casts some doubt on its warrant for an unequivocal “no” to Q, on which, see 1.4.4). To summarise, then, KT seem to be mounting an argument in

⁵ I am using the term ‘object’ here not in the sense of a physical object necessarily, but in the more general sense of the target or content of thought or perception. That is, as traditionally used in the dyad of ‘subject-object’.

which, if KT2 is confirmed, then the answer to Q is confidently no; if KT3 is confirmed, the answer is “nearly” no, and if KT1 and KT4 are true, we have *a fortiori* evidence to answer Q with a no.

1.3.2 De Brigard and Prinz (DP)

De Brigard and Prinz (2010) (henceforth referred to as DP) have argued for the “yes” answer—that attention is both necessary and sufficient for consciousness—so that attention and consciousness never come apart. I summarise their argument in the following four propositions. The first two propositions can be seen as the positive prongs of their argument, while the last two are defences against the empirical challenges raised against their view.⁶

DP1. *Attention and consciousness share striking functional similarities.*

DP2. *There is powerful behavioural evidence that “consciousness comes and goes with attention” (p. 53).*

DP3. *The evidence that attention is not sufficient for consciousness fails to establish that conclusion.*

DP4. *The evidence that attention is not necessary for consciousness fails to establish that conclusion.*

⁶ In fact, DP1 and DP2 can be taken as responses against KT1 and KT4, while DP3 refutes KT2 and DP4 refutes KT3.

Unlike KT above, the logical conclusion that their view is right cannot be drawn from the truth of any one of these propositions alone (1.4.1). The best they can hope for is that taken together, the truth of all (or many) of these propositions makes their position more likely. Even if all four propositions are shown to be true, neither the fact that *so far*, we have only observed attention and consciousness occurring together, and/or the fact that the evidence *so far* advanced against their concurrence failed to hold up can establish the truth of their position, since it will always be the case that just one piece of new evidence, just one incontrovertible new instance of a dissociation between attention and consciousness, will be enough for an answer of no to Q.

1.4 Analysis

A complete and exhaustive analysis of KT and DP would require a thesis all to itself, but here I am interested only in using them to introduce and illustrate certain issues that are highly relevant to my approach to Q in this treatise. I consider whether Q is an analytical or empirical question (or both); highlight the definitional ambiguity that characterises discourse on Q; raise the question of the functions of attention and consciousness; emphasise the lack of a comprehensive conceptual framework against which different approaches can be compared; point out the importance of working memory; suggest four questions that can improve the interpretation of empirical evidence; and provide an example of how easy it is to get such interpretations wrong.

1.4.1 Analytical and Empirical Questions

There are two important ways of framing and approaching Q that need to be considered. The first is the *analytical* approach:

Q_A: *Is attention analytically both necessary and sufficient for consciousness?*

Q_A requires an understanding of what kinds of things attention and consciousness are, what functions they perform, and from there, working out whether the very idea of attention is both necessary and sufficient for the very idea of consciousness. Our answers to Q_A depend heavily on the definitions of attention and consciousness we start with. KT define attention as volitional top-down attention, and are ambiguous as to their definition of consciousness. DP define attention as the gateway to working memory, and thereby to consciousness. In Chapters 2 and 3, I develop what I argue are more useful descriptions, with phenomenal consciousness as unified, temporal, situated first-person what-it-is-likeness, and attention as a suite of strategies for structuring cognition for further processing.

The other approach to Q is *empirical*:

Q_E: *Is attention empirically both necessary and sufficient for consciousness?*

Q_E may only be addressed once we have at least some of the groundwork necessary for an answer to Q_A. First, we cannot profitably analyse empirical data without having a clear idea about the concepts of attention and consciousness. This definitional ambiguity risks confusions I describe shortly (1.4.2). Certain answers to Q_A can make Q_E moot. With the right definitions of attention and consciousness—e.g., attention just is access consciousness, (9.2.2.3)—the answer to Q is trivially yes, since consciousness is defined *in terms of* attention, rendering the two inseparable. If the concept of “bachelor” is defined in terms of one’s marital status, then being a bachelor is trivially necessary and sufficient for being a male who never married. In defining attention as the gateway to working memory, and holding that all and only

content in working memory can become conscious, DP beg the question of Q_E (and Q) by their answer to Q_A .

While both papers make assumptions about the analytical nature of attention and consciousness and their relationship, the eight propositions I listed above (KT1-4, DP1-4) are all empirical propositions. They highlight an asymmetry in answering Q_E —*we can only be confident of a no answer, never of a yes answer*. KT2 and a KT3 without the “near-absence” qualification are the only propositions in both papers capable of delivering an unequivocal answer to Q_E (and therefore, to Q)—the rest could at best be suggestive, never conclusive. KT1, KT4, and DP1 cannot be taken as addressing Q_A , since they are matters of empirical investigation, at least as KT and DP treat them. Similarities and differences in *function* or *effects* do not equate to the necessity or sufficiency of one *process* for another (1.4.3 below). Showing a pattern of apparently consistent empirical (Q_E) *correlation* (DP2) is no guarantee that such correlation holds *universally* or is analytical. Undermining conclusions from the empirical evidence currently available (DP3, DP4) leaves open the possibility of new evidence, or even other arguments, that may still confirm those conclusions. Only KT2 and an unqualified version of KT3—the propositions that attention occurs empirically without consciousness and vice versa—if borne out, provide a definitive answer to Q_E . In summary, no amount of empirical evidence can supply a certain yes answer to Q_E (and therefore, to Q), while only one case of an empirical dissociation between attention and consciousness can supply a certain no answer. I argue in Chapters 7 and 8 that foveal phenomenal overflow is such a case.

On the other hand, Q_A does not exhibit this kind of asymmetry. Since the answer to Q_A depends only on the conceptual characters of attention and consciousness, it is equally possible to arrive with logical certainty at either a yes or no answer to Q_A . While Part I of this treatise attempts to clarify issues to do with the concepts of attention and consciousness and the relationships between them (Q_A), much of Part II of this treatise is devoted to consideration of such candidates for empirical dissociations (Q_E), although inevitably, the two questions overlap.

1.4.2 **Definitional Ambiguity**

Q_A evokes the importance of being clear on what we mean by attention and consciousness. Both KT and DP suffer from problems on both definitions. For KT, attention is defined as the process of selecting input

“defined by a circumscribed region in space (focal attention), by a particular feature (feature-based attention), or by an object (object-based attention) ... volitionally controlled forms of selective, endogenous attention” (p. 16).

Attention has been taken to mean both volitional and non-volitional, endogenous and exogenous (top-down and bottom-up) types of selection (3.2.3), and has even been taken to apply to things other than selection as such, such as access and detection (3.3.4.2). KT’s definition of attention then, can deliver an answer to Q only for a very narrow definition of attention. That is in itself a valuable exercise, but one liable to a very valid challenge that, for example, consciousness in the absence of top-down attention may still be consciousness in the presence of bottom-up attention.

Moving to consciousness, KT define it by its functions, which include:

“summarizing all information that pertains to the current state of the organism and its environment and ensuring this compact summary is accessible to the planning areas of the brain, and also detecting anomalies and errors, decision making, language, inferring the internal state of other animals, setting long-term goals, making recursive models and rational thought” (p. 17).

In other words, consciousness plays no role in selection—even top-down selection—but instead functions as the executive of the brain, coordinating its various subfunctions. They seem, therefore, to be identifying consciousness roughly with the executive function of the mind. An important point to note here is that while KT acknowledge the complexities involved in defining consciousness, and distinctions such as that between access and phenomenal consciousness (p. 16, see Block, 1995), their statement of the functions of consciousness quoted above describes *just* access consciousness, not phenomenal consciousness. All of the functions they describe in the quote above can be also performed implicitly, or subconsciously—without phenomenal consciousness (Hassin, 2013; Lamme, 2010, p. 210; Velmans, 2014). In fact, one could argue that what they have given is more appropriately considered a definition of something like a global workspace or working memory, rather than phenomenal consciousness as such. There is doubtless a close relationship between phenomenality and the cognitive executive, but it certainly cannot be simply *assumed* that they are identical. In Part II I consider the possibility held by leading theorists of working memory that all and only the content of the cognitive executive is phenomenally conscious, but I will in fact argue that there is good reason to think this possibility is wrong.

In later work KT do more directly show their interest in phenomenal consciousness as it relates to attention. In their response to Mole they speak about “consciously experiencing” objects or properties of objects, and compare this to “non-conscious priming” (Koch & Tsuchiya, 2008). Priming would normally be taken to be a case of access consciousness without phenomenal consciousness:

“The information-processing function of phenomenal consciousness in Schacter’s model is the ground of the concept of consciousness that I will mainly be contrasting with phenomenal consciousness, what I call ‘access-consciousness.’ A perceptual state is access-conscious, roughly speaking, if its content—what is

represented by the perceptual state—is processed via that information-processing function, that is, if its content gets to the Executive System, whereby it can be used to control reasoning and behavior” (Block, 1995, p. 229).

Elsewhere, they define what they mean by consciousness in the context of Q by equating consciousness with “the contents of conscious experience” (Koch & Tsuchiya, 2012). This strongly suggests phenomenal rather than access consciousness. So, while they have not made it clear in the 2007 paper that they are addressing phenomenal consciousness rather than access consciousness, it is clear from other papers that they are indeed interested in the relationship between attention and phenomenal consciousness. Perhaps their earlier work simply assumed that where access consciousness is, there also will phenomenal consciousness be? But this example from KT illustrates the fact that attention and consciousness come together so commonly as to tempt some to simply equate them to each other (Q_E), and that conceptually, attention and consciousness are easily construed as perfectly overlapping (Q_A). I argue later that there are very good reasons for rejecting both of these ideas.

DP fare little better, definitionally. Their definition of attention as the gateway to working memory, and thereby to consciousness runs the serious risk of begging the question of Q. On this model, content can never reach consciousness without passing through the sieve of attention first, so attention has been defined by the model as being necessary for consciousness. If all the content of working memory is conscious (this is an ambiguous point in DP), and if attention is the gateway *only* to working memory and nowhere else, then attention has been defined by the model as sufficient for consciousness. DP are quite aware of the danger of merely defining an answer to Q but have been taken by others not to have avoided the danger as well as they might have (3.2.4.2).

Contributing to the potential confusion is the fact that the large literature on Q, both philosophical and empirical, is by no means definitionally consistent. So, any paper

that attempts to synthesise an answer to Q from an array of other papers (as do KT, DP, and this treatise) must be constantly vigilant as to which definitions of attention and consciousness are in play, and ask whether or not the evidence adduced really justifies the conclusion reached without equivocation. Our challenge in Part I is going to be to come up with working definitions of attention and consciousness that respect the intimate conceptual and empirical connections between them while nonetheless providing principled grounds for addressing Q as an open empirical question, while consulting a definitionally unruly literature.

Ultimately, for the purposes of Q, one may define attention and consciousness in different ways, but one must be cognisant of the fact that different definitions will lead to different answers to Q (9.2.2). For example, different answers may eventuate for access consciousness versus phenomenal consciousness. And comparing either kind of consciousness against top-down attention may result in a different answer to that against bottom-up attention.

1.4.3 The Function Question

Q_A is in part a question not only about how we define attention and consciousness, but also about *functional roles*. The point of KT1 is that if attention fulfils a different functional role to that of consciousness (which DP1 denies), this is suggestive (though not conclusive) evidence of their dissociation and therefore, of the negation of Q. The point of KT4 is that if attention produces opposing *effects* to consciousness, that leads to the same conclusion.

“To the extent that one accepts that attention and consciousness have different functions, one must also accept that they cannot be the same process” (p.17).

But the assumption that a single process (or other kind of entity) cannot simultaneously play different functional roles or have opposing effects is by no means obvious. Most of us are well acquainted with the agony of a divided mind when faced with a moral dilemma. In chemistry and physiology generally, there are innumerable cases of complex systems pulling in opposing directions simultaneously, thus attaining an equilibrium of opposing effects, functions, or forces. For example, enzymes are biological catalysts, molecules in organic systems that enhance or impede chemical reactions. An enzyme like P450 is a complex protein that performs the same process—transferring a hydrogen atom from a substrate to an oxygen molecule—in a vast array of organic chemical reactions. The same process thus results in a wide variety of functions and effects, some of which, presumably, may eventually oppose each other.

In cognition itself, there is good reason to think that the same cognitive process may perform different functions in different circumstances, or given different inputs. The brain appears to be replete with the phenomenon of neural reuse (M. L. Anderson, 2010). There seems to be good evidence that it often happens that over evolutionary timescales, the same hardware and software operating in much the same way—the same processes—come to be used for novel functions, while either retaining or losing their original functions. So even if KT1 is established, it is not enough to answer Q, since attention and consciousness might simply be one and the same cognitive process that plays different functional roles and produces opposing effects in different circumstance (i.e., Q = yes), or they may be completely dissociable and independent processes that play different functional roles and produce opposing effects (Q = no). More information is needed, beyond KT1 or KT4.

What is missing is to determine the nature of the relationship between attention and consciousness. In 4.4, I consider what I believe to be the exhaustive list of possible kinds of relationship. Briefly, if attention and consciousness are *numerically identical*, then the answer to Q is certainly a yes, although even that would not necessarily be incompatible with the truth of KT1 or KT4, for the reasons described above. But this is not what KT are arguing, as I read them. If their relationship is one

of *causation*, e.g., attention causes consciousness,⁷ then attention will be necessary for consciousness but not sufficient. Further information about the circumstances of the causal relationship will be needed to elucidate whether attention is sufficient for consciousness. And if attention and consciousness are merely *correlated* without identity or causation, then of course the answer to Q is no, barring some weird coincidence that sees them always co-occurring without fail by pure chance.

The functional roles of attention and consciousness (and working memory) are quite pertinent to the arguments I offer in Part II, but they will play different roles to those they play for KT. Based on the definitions of attention and consciousness I develop in Chapters 2 and 3, I will argue that cognitive content can be phenomenal without being attended by the cognitive executive, and that this conclusion raises interesting questions about the purpose or functional role of phenomenal consciousness that are not easy to answer (9.5.2.2).

1.4.4 Comprehensive Framework

The discussion in KT is framed around four possible descriptions of a cognitive process: attended and conscious; attended but not conscious (KT2); conscious but not attended (KT3); and neither attended nor conscious. This matrix is much more comprehensive than the frameworks in many other discussions of Q, as I point out in 4.2.4. Yet KT's framework (and all other frameworks I have so far discovered) is logically incomplete and suboptimal for addressing Q. Q is framed by KT, DP, and others in the language of conditional logic—necessity and sufficiency. This is adequate to many discussions, but it does not capture the full gamut of possible relationships between two concepts such as attention and consciousness. Neither does it provide the kind of overall view that is required to answer Q for a whole

⁷ There are other possible relations of causation (4.4.1).

cognitive economy, rather than for a particular cognitive phenomenon such as gist perception. In Chapter 4, I develop a complete, precise, and optimal framework that does capture every possible scenario linking attention and consciousness in a cognitive economy.

KT has been criticised for arguing for “consciousness in the near-absence of attention” rather than “consciousness without attention.” I consider this kind of answer to Q to be much less interesting. Probing Q promises to shed light on the very nature of attention and consciousness through the ways in which they relate. To establish that consciousness can occur in the total absence of attention carries far weightier consequences than consciousness in the near-absence of attention. To merely assert that they can occur together in differing proportions is not nearly so significant as to assert that they can come apart completely from each other. The framework I develop in Chapter 4 subsumes graduated approaches such as that of KT, and more starkly, of Montemayor and Haladjian (2015), and allows such approaches to be compared with others that do not share this trait.

1.4.5 Working Memory

We saw above that at least in this 2007 paper, KT’s definition of consciousness is very much like Block’s access consciousness, Baars’ global workspace, or the central executive of working memory models. We shall see later that many authors simply take this cognitive executive to be synonymous with either attention, or consciousness, or in some cases, both. In this treatise, I argue against this assumption. KT and DP agree in ascribing different roles for attention and consciousness, even if they arrive at opposing answers to Q. For KT, the different roles suggest that attention and consciousness can dissociate (Q = no). For DP, attention is the gateway to working memory, and consciousness arises only from the workings of working memory, so there can be no consciousness without attention, and attention will always produce consciousness (Q = yes). These discussions imply

that the concept of working memory (whatever that may be) is likely to be most helpful in elucidating the nature of the relationship between attention and consciousness, and thereby answering Q. To that end, I survey the principal working memory literature in Chapter 5 and use it in Part II to develop my arguments.

One of the ways working memory will prove useful is in the definition of a particular kind of attention, Executive Attention, that will play an important role in my arguments. In 3.4.4, I argue that attention as a cognitive strategy is virtually ubiquitous, which therefore makes Q trivial. Not only are all conscious contents attended in this broader sense of attention, virtually *all* cognitive contents—even unconscious contents—are attended. But attention implemented by a cognitive executive is far more circumscribed. How can such attention be delineated in a principled manner? Models of working memory provide exactly what is needed.

1.4.6 Four Questions

Both KT and DP (and most other authors) invoke numerous empirical studies to support their arguments. But much of the work thus invoked was not specifically designed to directly address Q. Even work that is so designed can suffer from a degree of interpretive confusion that clouds its application to Q. To avoid this kind of confusion, and to more clearly determine just what can validly be drawn from any given experiment with respect to Q, I propose four salient questions that can be asked about any experimental paradigm, which I then employ where helpful throughout this treatise, and capitalise to indicate that I am referring to these specific questions.

A. Target

What is the content of consciousness and what is being attended?

Do attention and consciousness have the same target?

B. Timing

When does the attention occur (or not), and

When does the consciousness occur (or not)?

C. Variety

What type or variety of attention is occurring?

What type or variety of consciousness is occurring?

D. Consequences

What exactly changes due to the experimental manipulation?

Target Question

The *Target* question asks whether the attention and the consciousness pertain to exactly the same Target, and prevents erroneous conclusions about the correlation between attention and consciousness. A very uninteresting version of Q would be whether it is possible to attend to X while being unconscious of Y, e.g., can I attend to a Beethoven symphony while being unconscious of the taste of a juicy grape in my mouth? Clearly, what we really want to know is whether I can attend to the music without being conscious of that same music, or vice versa.

Timing Question

The *Timing* question ensures that the particular instances of attention and consciousness in question are simultaneous. If attention begins at T and ceases at T+1, and then, consciousness—even of the same Target—begins at T+2, then there is an important sense in which attention and consciousness did not co-occur—the temporal sense—which is highly relevant to Q.

Variety Question

In Chapters 2 and 3 I describe some of the distinctions people have made between different types or varieties of consciousness and attention respectively. Keeping track of which variety of each an author is talking about can play an important role in preventing a kind of equivocation that can confuse matters. This is the essence of the *Variety* question: which specific kinds of attention and consciousness are we talking about, and do we jump illegitimately from one kind to another within the same experiment and/or argument? For example, it is unhelpful in relation to Q to determine that I am spatially attending to the locus of a speckled hen, but unconscious of some of the spots on that hen (which would correspond to feature attention, not spatial).

Consequences Question

The relevance of some of experiments to Q depends on certain putative consequences of manipulating attention or consciousness. Thus, it will be important to consider what exactly these consequences are, and whether or not they establish the things that authors need from them for their arguments.

1.4.7 Pitfalls of Empirical Interpretation

Just like prose or poetry, empirical data necessarily require *interpretation*. This is a subtle art with many pitfalls that can trap even the most experienced and capable practitioner. In addressing Q, we will need to be as vigilant and self-critical as possible. An example from KT illustrates this danger, and the usefulness of the four questions above, but there are many examples in DP and many other papers.

In support of KT4—attention and consciousness sometimes have opposing effects—KT cite Olivers & Nieuwenhuis (2005) who found that distracting a subject seems counterintuitively to improve their attention by reducing the *attentional blink*

effect.⁸ Subjects were asked to identify two numerical digits (denoted T1 & T2) inserted into a rapidly presented visual series of letters. Attentional blink means that within the right range of duration for the gap between the two digits, performance on identifying T2 should be poor, and that is exactly what happened with the subjects who performed the experiment in the standard way. However, two other groups of subjects performed the same task while concurrently distracted with either a free association task (thinking about a holiday or a shopping list) or with listening to rhythmic music. Oddly enough, the distracted groups were *better* at identifying T2 within the attentional blink timespan than the standard group. KT take this to be a case where decreasing attention through the distraction, actually increases conscious awareness, as reflected in the improved accuracy of identification of T2.

This provides an opportunity to illustrate the benefits of my Four Questions. The *Targets* of interest are the series of presented letters and numbers, and more specifically, T1 and T2, while the *Variety* of attention is top-down selective attention in the standard case, but bottom-up attention in the distractor cases. In all cases, spatial attention is focused on the location of the presented letters and numbers. It is continuously so focused in the standard case, but continuously distracted in the distractor cases (*Timing*)—not spatially, but through the recruitment of cognitive resources to largely non-visual tasks. The *Consequence* of the distraction of attention is that recognition of T2 within the usual attentional blink window improves.

The case KT can make, then, is that during that attentional blink window, cognitive resources are directed away from the task of processing and recognising the presented letters and numbers generally, and from identifying T1 and T2 specifically. This *drop* in attention paradoxically results in better recognition of T2 (T1 recognition is not significantly affected)—conscious awareness of both, but

⁸ Attentional blink (Raymond, Shapiro, & Arnell, 1992) is the phenomenon where a second stimulus that closely follows on a first stimulus (within 500msec) is not consciously perceived, most likely because the limited attentional resources are still occupied with the first target stimulus, and therefore have insufficient resources left over to recognise the second stimulus.

especially T2 in the attentional blink window, is *increased*. From this they conclude that attention and consciousness have opposing effects.

The argument KT wish to make is not one that can be made from this experiment. KT phrase it thus: “reducing attention can enhance awareness” (Box 2, p. 20). But an argument of the form: “*reducing X increases Y*” is not an argument for attention and consciousness having opposing effects, but for attention and consciousness varying in inverse proportion to one another. What KT need to argue to establish opposing effects is something like this: “*reducing X increases Z, but reducing Y decreases Z.*” In other words, reducing attention increases the occurrence of a Consequence for a particular Target, while reducing consciousness decreases the same Consequence for the same Target. That is not the structure of Olivers and Nieuwenhuis’ experiment.

Perhaps what KT want to say is that taking T2 as the Target of both attention and consciousness, in the Olivers and Nieuwenhuis experiment, the Consequence of the manipulation was that consciousness went up when attention went down. But this is still not KT4, since this way of phrasing what happened holds the modulations in both attention and consciousness as Consequences, not causes, and KT4 requires both attention and consciousness to be causes with differing effects. This is a poor way of interpreting the experiment and I doubt very much that this is what KT are doing. In fact, any interpretation that holds consciousness as a Consequence rather than a cause cannot support KT4.

Another problem is that Olivers and Nieuwenhuis themselves offer three possible interpretations of the reasons for this effect (p. 268), none of which are friendly to KT4. First, they wonder whether increased arousal may have resulted in more attentional resources coming online generally, thus increasing the attentional resources proportionally for that part of attention devoted to identifying T2. However, in their reward group, when subjects had good financial motivation for being aroused, there was no improvement in T2 identification. If this interpretation is correct, then the experiment becomes a case of increased attention *and* increased consciousness, which are not opposite effects. Second, music or thinking of

pleasurable things may produce a positive affective state, which has been shown to improve many cognitive functions. To test this hypothesis, one would need a distraction that was affectively neutral. Once again, this would mean increased attentional resources for T2, and the experiment becomes a case of increased attention *and* increased consciousness, which are not opposite effects. And third, the distractor task may have actually had the effect of widening the *duration* of the spotlight of attention, thus encompassing both T1 & T2 in the one spotlight. This explanation also counts against KT4, since it implies that the improved consciousness of T2 is consequent to greater attention being paid to T2.

This is one of the pitfalls inherent in utilising empirical studies to answer a question like Q that were not specifically designed for that task. Olivers and Nieuwenhuis themselves draw no conclusions in their paper about the relationship between attention and *consciousness*. The words “conscious” and “aware”⁹ do not figure significantly at all, anywhere in their paper. Their conclusions, instead, focus on highlighting the complex structure of attention itself—withdrawing attention in one way seems to improve attention in another way. In my arguments in Part II, I strive to avoid these kinds of pitfalls.

1.5 Outline of Thesis

It has been said that the definition of insanity is to keep doing the same thing over and over and expect a different result.¹⁰ The work on Q is a body that is large, impressive, and creative. Out of respect for its quality, I have tried in this treatise *not*

⁹ The word “aware” does appear once, but only in a methodological sense: “None of the participants were *aware* of conditions other than the one they were placed in” (267).

¹⁰ Attributed to Albert Einstein.

to do the same thing over and over that has been done by others far better than I ever could. Instead, I have focused in Part I on a careful reassessment of the very foundations of the discourse. It turns out there are many different ways Q can be posed, and the literature is not always as clear as it might be on which particular version of Q is being addressed. In Part II, I choose a particular version of Q and probe it methodically, in light of the foundations laid in Part I, gradually narrowing down the investigation, until I construct an abductive case for a “no” answer to a specific framing of Q. The last chapter includes not only the results of this investigation, but some lamentably brief indications of possibly fruitful directions for future investigation. Below I outline the structure of this treatise chapter by chapter, highlighting where possible, what is unique to my approach.¹¹

1.5.1 Part I

Chapter 2 addresses the definition of phenomenal consciousness as I use it in Part II. I define phenomenal consciousness by identifying four core characteristics: *what it is like-ness*; *a situated first-person perspective*; *phenomenal unity*; and *temporality*, all four of which are necessary together to identify an instance of phenomenal consciousness. I further draw the distinction between the *content* of phenomenal experience and *phenomenality* as such.

¹¹ A few months before the submission date for this thesis, the *Philosophical Transactions of the Royal Society B: Biological Sciences* published a Theme Issue (2018, volume 373, issue 1755) on “Perceptual consciousness and cognitive access,” compiled and edited by Peter Fazekas and Morten Storm Overgaard. The seventeen cross-disciplinary articles cover much of the same ground as this treatise. My work was largely completed prior to its publication, and some of my most important arguments were presented at the IACAP (21-23 June) and ASSC22 (26-30 June) conferences in Poland. Nonetheless, I have endeavoured to engage with as much of that issue as is relevant, and as space and time allowed. For example, the challenges raised by Gross (2018) against arguments for phenomenal overflow from capacity limitations of working memory are directly relevant to my arguments in Chapter 8, but would require a whole thesis to themselves to do them justice.

Chapter 3 turns to the definition of attention, an area where significantly less work has been done, compared to that on consciousness. After surveying and analysing other approaches and some common distinctions among types of attention, I take a descriptive approach, surveying the literature and collating and taxonomizing the various ways attention has been defined. I conclude the chapter by distilling a definition of attention that lies at the heart of the Operations definitions in the taxonomy: attention is *a suite of strategies for structuring cognition for further processing*. I conclude by drawing an important distinction between the ubiquitous Liberal Attention and the much more circumscribed Executive Attention, the latter being the definition I employ in Part II.

In Chapter 4 I turn to the nature of the relationship between attention and consciousness. Starting from scratch, I build a uniquely comprehensive Set Theoretic Framework that captures all the possible logical relationships between two entities within a cognitive economy and compare it to existing frameworks in the literature. An answer to Q may obtain over different domains, which I outline in terms of weaker and stronger readings. I then situate nearly fifty authors against this framework and draw some general reflections from this review of the literature. I conclude the chapter by describing four possible kinds of relationship that might underlie the pattern of occurrence described by the Set Theoretic Framework. My answers to Q in Chapter 9 will be framed in terms of these constructs.

In Chapter 5 I describe the leading models and metaphors of working memory, a concept that frequently appears in the discourse on Q. I use these models to pose a number of pertinent Questions that illustrate how huge Q is. While I touch upon them all at some stage, I will have the space to address only one—the Capacity Question in Chapter 8—in any substantial detail. They are certainly not original questions, or unique to Q, but they help guide investigation of Q. Working memory serves many purposes in this treatise, among them, providing a principled way to delineate the boundaries of Executive Attention from Liberal Attention and probing the capacity limitations of attention and consciousness. I conclude with some observations about the dual-nature of working memory: storage and manipulation.

1.5.2 Part II

In Chapter 6 I begin broadly by searching for evidence for the two kinds of cognitions that will settle Q: attention without consciousness, and consciousness without attention. The first proves relatively easy to establish, but the second is much harder. To facilitate the search, I invoke my definitions of attention and consciousness from Part I to derive five classes of ways that there might be consciousness without attention in cognition, whether in specific cognitions (“local”), or in a cognitive economy as a whole (“global”). None of these turns out to carry conclusive evidence of consciousness without attention.

The search narrows in Chapter 7 as I turn to a particular form of local consciousness without attention: phenomenal overflow. But I bypass the usual objections to overflow, which all relate to peripheral vision, and focus on phenomenal overflow in *foveal* vision. Here, I argue, there is good reason to hold as veridical a subject’s confidence that he consciously sees every speckle on a speckled hen clearly and distinctly, even though that subject cannot further process (Executively Attend) every speckle. I present arguments from the temporal nature of conscious experience and the immediacy of experience, in the sense that one’s experience of perceptual content itself (not the object out there in the world) can be unmediated by attention. I conclude the chapter by considering whether no-report paradigms might offer a way of bypassing the epistemic obstacle of relying on report by the subject (which requires Executive Attention) to identify conscious content. Unfortunately, not only do extant no-report paradigms fail to identify consciousness without attention—there is good reason to think they can *never* succeed.

In Chapter 8 I build another abductive argument for foveal phenomenal overflow, this time by invoking the Capacity Question from Chapter 5 and applying it to attention, consciousness, and working memory. Triangulating these three capacities strengthens the case by showing that there is no reason why the three capacities—or even consciousness and any one of the other two—should be identical. In fact, given how differently they behave, it would be an incredible coincidence if their

limitations were identical. I express this in a Witches' Hat Model of cognitive content and processing, on which phenomenal overflow is almost inevitable, and then compare this model to three prominent alternative accounts: the illusion of richness; expanded attention; and inchoateness. I conclude that the Witches' Hat Model with phenomenal overflow is the account that best explains all the data.

Finally, in Chapter 9, I bring all these arguments together to summarise my answer to Q in terms of the foundational concepts built up in Part I: the Set Theoretic Framework; the kinds of relationship that might underlie these patterns of occurrence; and the readings of the domains over which they might obtain. I provide answers to four versions of Q, focusing mostly on that concerning Executive Attention and phenomenal consciousness, the chief focus of Part II. I conclude the treatise with some reflections on the implications and consequences of my arguments for the discourse on Q, and on other questions: the nature of attention and consciousness; theories of consciousness; other cognitive concepts (intelligence and symbolic cognition); and some practical ethical issues involving artificial intelligence, animals, infants, and the ascription of moral value and moral responsibility.

PART I

(Re-)Laying the Foundations.

2 Phenomenal Consciousness

2.1 Phenomenal Consciousness and Q

“Consciousness is a word worn smooth by a million tongues”
(Miller, 1962, p. 25).

There are many varieties of Q that could be posed—of lesser or greater interest and utility—depending on what one means, exactly, by both attention and consciousness. In this treatise, I am interested in a particular variation on Q, one in which “consciousness” signifies *phenomenal consciousness*. In this chapter, I develop a particular concept of phenomenal consciousness, not because I think it the only “right” concept, but because I think it the most interesting and useful concept of consciousness that can be addressed in Q. I develop a particular concept of attention in Chapter 3, although it will require significantly more effort, since attention has been significantly less studied than consciousness, philosophically at least.

I begin my analysis of phenomenal consciousness by briefly surveying the relevant history of the concept and the term “consciousness” before outlining in some detail the particular concept of phenomenal consciousness I intend to use, defined by its intrinsic features. I conclude with a discussion of a very pertinent distinction between the content of phenomenal experience and phenomenality as such, in itself.

2.2 A Brief History of “Consciousness”

While the philosophical term *consciousness* as we define it today was introduced only as late as the eighteenth century,¹² the concept of being conscious likely predates recorded history, being as it is, the central defining feature of human experience,¹³ and was not limited to human beings—“spirits” were ascribed to animals and plants and even natural features such as rivers and stars.¹⁴ Terms roughly equivalent to the English “soul,” “spirit,” “mind,” “intellect,” and so on abound in the ancient languages of the earliest civilisations such as Mesopotamia, Egypt, and Greece. Seager (2007, pp. 12–13) posits that the earliest civilisations adopted one of two contrasting approaches to understanding mind or consciousness: *emergence*, by which consciousness is constructed out of non-conscious components; and *panpsychism*, by which consciousness is an elemental or *sui generis* property of reality. It is with the Presocratics (and, at roughly the same time, in Indian philosophy (Dreyfus & Thompson, 2007)) that we see the first

¹² Cudworth (1678), see also Lähteenmäki (2010).

¹³ Controversial ideas such as the bicameral mind notwithstanding (Jaynes, 1976).

¹⁴ For example, it has been suggested that some paleolithic hominins saw elephants “as a ‘sister-species’, resembling humans in physical, social, behavioral and perceptual aspects” (Lev & Barkai, 2016, p. 240).

recorded attempts to rationally and systematically evaluate these broad approaches and elaborate an explanation of mind. The current debate over attention and consciousness is the direct descendent, in some ways perhaps the continuation, of these attempts. Of course, the discourse has progressed greatly, as has the context in which the discourse is conducted—our overall understanding of the world in which we live. Empedocles grappled with how mind emerges from the four elements, Democritus found himself unable to explain how experience emerges from atoms, and Epicurus later grounded our volition in the swerve of these atoms. While we now know so much more about atoms, and about the many layers or levels of organisation atoms can form, all the way up to functioning biological organisms, some may argue that we are still approximately as far from explaining consciousness as they were.

The difficulties inherent in grappling with so mysterious a concept as phenomenal consciousness has kept many brilliant modern minds away from the issue, particularly scientists.¹⁵

“Consciousness used to be a controversial topic of study. Not only during the heyday of behaviorism, but also during the rise of cognitive science in the 1970s and 1980s, only a few senior scientists (such as Gerald Edelman and Francis Crick), who had first achieved success and job security in completely unrelated fields, felt free to attack this final, big question. But how things have changed in the past 20 years!” (Block et al., 2014).

According to Baars (1988, p. 34), the return to respectability of the scientific study of consciousness began with the work of Donald E. Broadbent in the 1950’s. The work of Broadbent and others on “rapid shadowing”—listening to a stream of

¹⁵ Scientists, perhaps, but not psychoanalysts, artists and novelists, as Wertheim (2016) points out.

speech and immediately repeating what you hear—not only provided an experimental paradigm for probing the structure and limitations of consciousness, but also illustrates the close relationship between attention and consciousness. The subject voluntarily learns to selectively attend to the stream of speech to the exclusion of all other sensory input, and is (apparently) conscious of only that stream. Interestingly, Moray (1959) found that sensory input in the other ear—for example, the subject’s name—could be processed unconsciously and “break into” the subject’s own stream of speech. In spite of this close connection between attention and consciousness at the very beginnings of the modern revival of interest in consciousness, the explicit *relationship* between the two featured little in scholarly discourse until the late 1980s (4.1.1).

2.3 The Ontology of Consciousness

With such a long history and universal access to it, the term consciousness has come to mean different things to different people in different contexts, so a clarification of the kind of consciousness I am interested in in this thesis is in order. I am interested in *phenomenal consciousness*, in the sense of Block’s (1995) phenomenal consciousness, or in the sense made famous by Thomas Nagel’s (1974) paper *What Is It Like To Be A Bat?* There is “something it is like” to be you, right now, reading this treatise. It is very likely that there is nothing it is like to be the codex or the computer upon which you are reading this treatise (although far more of the content of the treatise resides in both than in your brain).

I am not interested here in *physiological* or *clinical* consciousness, such as that measured by the Glasgow Coma Scale in hospital emergency departments, nor in consciousness as Block’s (1995) *access consciousness*, which is the availability of

information for further processing,¹⁶ nor in *knowledge* or *insight*, in the sense that Buddhists speak of the consciousness of enlightenment. Further, I am not here interested in the *universal* (metaphysical category) “consciousness,” but in actual instantiations of consciousness—particular conscious experiences like yours and mine.

A great deal of the interest in understanding consciousness has taken the form of searching for a *theory of consciousness* (9.5.2), and the vast and varied array of viable contenders is testimony to our befuddlement. Surprisingly, though, there has been relatively little said about the question of what *ontological category* phenomenal consciousness might belong to. If we employ a straightforward ontological taxonomy (Lowe, 2002, p. 16),¹⁷ the *concept* of “consciousness” is easy to classify—it belongs to the metaphysical category of *universals*. But whereas the *particulars* of other universals are also easy to classify, the particular of the universal “consciousness”—particular conscious experiences—is not so easy. An individual dog is a particular concrete *substance*. The Grand Final of a sports competition is a particular concrete *event*. The redness of this apple is a *property* that instantiates the universal property of redness, and its connection to this apple is a *relation*. But to which category does your current conscious experience belong? Objects? Events? Properties? Relations? Or something else? This fascinating question is one I address elsewhere, where I explore the possibility that the uniqueness of phenomenal conscious may warrant its having its own ontological category.¹⁸

One approach to Q would be to ascertain the ontological categories of both attention and consciousness, and thereby ascertain the kinds of relationship they might be capable of having. For example, if attention is an *event* and consciousness is a *property*, then Q could be posed as a question about whether consciousness is a

¹⁶ Although we will later see that Block’s access consciousness bears a very close affinity to what I shall define as Executive Attention (7.2.2).

¹⁷ There are, of course, many other ontological schemes, some of which may be better suited to categorising consciousness. However, I use Lowe’s scheme here purely for illustrative purposes.

¹⁸ Paper currently under preparation.

property always and only of attentional events. However, the ontological categorisation of consciousness is a very difficult endeavour indeed—much harder than attention—and one which I will therefore eschew in this treatise in favour of just being content to identify its defining features so that we have a principled way of ascertaining whether or not it is present in cognitions of interest.

2.4 Defining Phenomenal Consciousness

A large part of answering Q is going to involve identifying the presence or absence of phenomenal consciousness and attention. This is how the framework I develop in Chapter 4 works. For the purposes of Q, we need not worry about allocating consciousness to an ontological category, we do not need an all-explaining theory of consciousness, nor must we enter into the finer points of defining it. We just need a set of signature features that allow us to ascertain whether or not consciousness is present or absent in any given circumstance—a *natural kind* approach (Shea & Bayne, 2010).

What makes phenomenal consciousness what it is? What are its distinguishing features? One approach might be to consider Chalmers' famous (and controversial) philosophical zombie thought experiment (Chalmers, 1996). We can sift out the features that define phenomenal consciousness *qua* phenomenal consciousness by subtracting everything that the philosophical zombie can do from everything that a conscious person can do. Whatever is left over—that is phenomenal consciousness.¹⁹ Thus, for example, memory is not consciousness (even if there are

¹⁹ Another way to arrive at the core features of phenomenality might be to eliminate conscious content one by one till nothing is left. I explore something very much like that in 6.3.3, but we must beware of begging the question of whether it is possible to phenomenally conscious without being phenomenally conscious-of anything at all.

conscious memories) since the zombie is capable of memory. Whether via this route or other routes, I propose four tentative features²⁰ that together might form a *nomological cluster*—a set of evidential properties that occur together in nature (Shea & Bayne, 2010, p. 471)—that identify consciousness: “what it is like-ness;” the situated first-person perspective; unity; and temporality.

2.4.1 What It Is Like-ness

The first and most obvious difference between a philosophical zombie and a conscious person is that elusive idea of phenomenality—qualia²¹—“what it is like” to be you. Most authors surrender to this idea and do not even try to define it further.²² It is usually defined by appeal to your own conscious experience—if you are reading this, then you will just *know* what it is like to be you, reading this. No one else can know that. In fact, the Problem of Other Minds (Hyslop, 2018) suggests that no one can know with certainty that any other person is not a philosophical zombie. But as Descartes (1637) famously argued, the one thing I can be certain of is that I am a thinking (read, “conscious”) thing, else how could I be aware that I am asking the question of what exists?

I further assume an ancient Indian philosophical idea found in Dignāga and others: what it is like-ness can be purely *reflexive*—it is inherent in my experience that it is *I*

²⁰ Block (1993, pp. 313–317) lists eight features and van Gulick (2017, sec. 4) lists seven.

²¹ For a succinct review of the terms “what it is like,” “experience,” and “qualia,” see Bayne and Montague (2011, pp. 8–11). “Qualia” tends to refer to perceptual phenomenal experience, but of course, some hold that there are non-perceptual phenomenal experiences, such as thinking of an abstract concept like the number 937. This controversy, sometimes called the controversy over *cognitive phenomenology*, is not a controversy I can address here. For my purposes, I am just thinking about *any* phenomenally conscious experience whatsoever.

²² An excellent analysis of the concept, though, is Stoljar (2016).

having the experience—without being *reflective*—without my knowing or being *aware* in a second-order way that it is I having the experience (Ram-Prasad, 2007, p. 67).

2.4.2 Situated First-Person Perspective

The second feature is closely related to the first, but worth emphasising in its own right. This is subjectivity, the *first-person perspective*. Consciousness necessarily requires a subject as well as an object. For there to be an experience, there must be an experiencer of some kind.²³ And the experience is experienced from the unique perspective (spatial or otherwise) of that experiencer, and no other. Even in situations where we empathetically experience things from the point of view of another subject, such as when you wince at another's pain, you are not truly experiencing that same instantiation of pain that the other subject is experiencing, you are merely *simulating* the other's experience within your own experience and from your own perspective, with varying degrees of accuracy. This is true on any theory of empathy one might prefer (Zahavi, 2008), and the consequences of empathy may often work against the person for whom you have empathy (Bloom, 2017), emphasising that this is still your unique perspective, and not truly another's. The zombie, on the other hand, has no such experiencer to occupy this first-person perspective.

²³ Galen Strawson (2017), who argues for a variety of physicalist monism, still concedes that "Certainly all experience requires a subject of experience, a subject that lasts at least as long as the experience lasts" (p. 375). See also Strawson (1994, pp. 129–134), where he explores in detail what the necessity for an experiencer might mean, and (2016, pp. 92–93), where he explicitly argues for an experiencer, although not one that is ontologically over and above the experience. See also Duncan (2017) for a defense of the integral necessity of a self in phenomenal experience.

What is more, this unique first-person perspective is uniquely *situated*. This can mean different things: a subject's point of view can be bodily, rational, phenomenological, or a co-incidence of any or all of them (Rovane, 2012, p. 21). In spatial terms, what it is like to be me is not experienced from the vantage point of the whole universe, nor from a Paris café (sadly), nor from where you happen to be located in space, but from where I happen to be located in space at this moment in time (Merker, 2007, pp. 72–73). Autobiographical memories retain this situationality, even when they are remembered from a third person perspective (McCarroll, 2018). Even out of body experiences (Agrillo, 2011; Blackmore, 1991)—the sense of floating above one's own unconscious body—and ego dissolution (Letheby & Gerrans, 2017)—the blurring of the boundary between the self and the universe—retain this situationality. The locus may change, or it may enlarge, but there remains a locus. But situationality involves more than just spatial location. It is embodied, embedded, enactive, and extended (Menary, 2010)—its nature entwines inextricably with that of its environment. Conscious experience is also situated in time (2.4.4) and situated within a broader and complex context of non-spatial factors such as memories, emotions, goals, affordances, etc.

A related question is how we individuate conscious entities. For example, on an ancient Neoplatonic view, or on some versions of a modern panpsychist view, your personal experience of consciousness is not just the instantiation of “consciousness” (as a member of the metaphysical category of universals), but an outcrop of a greater “world consciousness.” What makes your consciousness distinct from other consciousnesses? Is there really just one big consciousness, or many? The situated first-person perspective inherent in our conscious experience provides a clear way to delineate my consciousness from yours, or from a world consciousness. It also provides a remarkably resilient way to delineate my conscious experience in the face of extended views of cognition which see the cognitive economy as extending beyond the integument of the human body (9.5.3.2). Whatever it might be that is carrying out the processes underlying my cognitions, I continue to phenomenally experience them from my own, single situated first-person perspective.

2.4.3 Phenomenal Unity

The third feature of consciousness, one related to the two above, is its remarkable *unity* (Bayne, 2010; Bayne & Chalmers, 2003; Brook & Raymont, 2014). Phenomenal unity seems to be intimately tied to what it is like-ness and the situated first-person perspective. It seems odd to think that at any moment in time, there could be more than one thing that it is like to be you. If there were two distinct things it is like to be you, then there would be two of you, not one.²⁴ It would also be odd to conceive of two perspectives somehow being a single perspective. Spatially, we think of a perspective as the outlook from a dimensionless point in space, one that is utterly distinct from the perspectives from other points. To move perspective to a different point is simply to take the perspective of the other single point. And even the binocular perspective we experience in normal vision is experienced as essentially a richer, congruent, single perspective—a perspective from a single point in space somewhere between the two eyes—rather than two discrete perspectives from the two independent loci of two eyes.

Phenomenal unity in no way limits the richness of phenomenal experience. The distinction between the content of consciousness and the phenomenality as such (2.5) is helpful here. Although the *content* of our synchronic experience is usually multiple and varied—within sensory modalities, across them, and beyond—*phenomenality as such* displays a remarkably resilient singularity of nature. I favour a *no experiential parts* view of consciousness unity (Brook & Raymont, 2014; Lee, 2014; Raymont & Brook, 2009) on which the plural *contents* of a synchronic phenomenal experience populate a solitary unified *phenomenal experience* (Tye,

²⁴ This point does not require commitment to realism about the self—only to an “experiencer” of some kind—see footnote 23 above. Nor does this imply you cease to be you when you are in a dreamless sleep. I only wish to make the point that so far as we can tell, phenomenal experience in humans cannot be divided in the way that, say, information in a computer can, or attention in a human can (on which, see 8.3.2).

2003). This is in contrast to an *experiential parts* view where many discrete synchronic experiences (hearing Beethoven, tasting a grape) are somehow glued together into a single composite experience. Bayne (2010) presents a subtle and fascinating account that combines holism and mereology.²⁵ All of the different contents of experience are unified and phenomenally interdependent²⁶ (Dainton, 2000, p. 191)—the different contents of a synchronic experience influence each other’s phenomenal character, as is highlighted in some kinds of visual illusion (B. L. Anderson & Winawer, 2005), but is happening all the time in normal cognition. This phenomenal interdependence emphasises the unity and indivisibility of synchronic phenomenal experience—there is always only ever *one* thing it is like to be you at any moment in time, albeit with multiple contents. By contrast, there is no such simple unity to be found in the philosophical zombie, in whom there is only multiplicity of interacting yet discrete content.

2.4.4 Temporality

The last feature is not one that arises out of the difference between the conscious person and the zombie, but out of the three features above. The kind of unified, situated first-personal, what it is like-ness we experience, as far as we can tell, is necessarily *temporal*. There is a situated perspective in time as well as in space. Experience is experience through time. There is nothing it is like to be you at just a dimensionless point in time. Nor can we can imagine such consciousness in the absence of time. There may be some other form of consciousness that does not need time, but the consciousness we experience and know seems to depend on time passing. Our common experience is not consciousness of a single dimensionless point in time, but of a specious present (James, 1890, p. 609), a span of a few

²⁵ I explore his account in my Masters thesis (Kaldas, 2015).

²⁶ Compare the idea of integrative binding (de Vignemont, 2014, p. 130)

seconds around that point in time,²⁷ and there are many temporal idiosyncrasies to our ongoing phenomenal experience (Herzog, Kammer, & Scharnowski, 2016). This is no doubt because of the way that our brains model reality for us to experience, but it remains true that we cannot imagine experiencing reality in any other way.

Indeed, time pops up in all sorts of ways in our exploration of consciousness. Cognitive processes happen over time, and there are interesting questions to ask about how phenomenal experience relates to the processes that produce it and its contents. Whether or not specific content is both attended and conscious requires us to clarify whether that content was both attended and conscious *at the same time*, or whether it was, say, first attended but unconscious, and then became conscious but no longer attended (the Timing Question, 1.4.6). And, when we identified consciousness as being unitary, we had to speak about consciousness at a moment in time. Whether consciousness is unitary *over* time is another question entirely. If it is, then it seems plausible that it is unified in a different way to synchronic consciousness. My experience right now and my experience tomorrow at midday do not seem connected in the same way that my experience of the taste of a grape and the sound of a Beethoven sonata at the same moment in time are connected. But this too is an issue too involved to do any justice to it here. For my present purposes, it is synchronic phenomenal unity that is the feature of interest in identifying consciousness.

2.4.5 Using the Features of Consciousness

All four of these features must be present to confirm the presence of consciousness. The absence of even one of the four is enough to discount the presence of consciousness. Temporality, of course, is a feature of all physical reality, and can

²⁷ See Appendix 1, *Synchronic v Diachronic* for more on the duration of the specious present.

therefore occur without the other three. But it seems unlikely that any of the first three can occur in isolation from the others.

What it is like-ness is, I believe, unique to consciousness and necessary for consciousness. It is the *sine qua non* feature, an *a priori* defining feature, much as being unmarried is a defining feature of bachelorhood.²⁸ The rest are not *a priori* features, but *a posteriori* features. Perhaps the world might have been such that they were not defining features of consciousness, but in our reality, they seem always to characterise consciousness, and it is very difficult indeed to conceive of conscious experience without them. All these features have been challenged. The least controversial is the first one, what it is like-ness, which has been the commonest way of describing consciousness in recent decades. I will not enter into the controversies surrounding these features here, but I ask the reader to grant that these are plausible, perhaps even somewhat likely features of consciousness for the sake of being able to identify the presence or absence of consciousness as we pursue our quest to answer Q.

2.5 The *Content-Phenomenality* Distinction

There is one more distinction to be drawn that is immensely relevant to my arguments: that between the content of conscious experience, and the phenomenality as such of conscious experience.²⁹ *Content* is comprised of the objects—in the widest sense—of phenomenal experience. Perceptions of the smell, shape, and feel of an apple are all conscious content, as are imagined white unicorns, arithmetical sums, ideas for a novel, and feelings of sadness. Content can of course

²⁸ Bachelors of Arts notwithstanding, of course.

²⁹ This distinction is found for example, Tye (2003) and Bayne (2005, 2010).

be either phenomenal or not, but where it is phenomenal, I generally use the term, “phenomenal content.” I use the term “content” to avoid getting into debates over the nature of the information in cognition. My argument here is neutral with regard to whether conscious content is representational or not, perceptual or not, etc.

On the other hand, *phenomenality as such* (hereafter, just “*phenomenality*”)—the phenomenal nature of experience in itself—is conceptually distinct from the objects or content of that experience. Phenomenality is chiefly captured by the ideas of what it is like-ness and a situated first-person perspective, and is much harder to characterise than content. It is not itself content—not itself an object of experience—but experience itself, subjectivity itself.

“the one necessary constituent of consciousness that can never be an object of consciousness is that very vantage point itself, namely, the origin of the coordinate system of the simulation space. It cannot be an object of consciousness any more than an eye can see itself” (Merker, 2007, p. 72).

Content and phenomenality have different attributes. For example, Bayne et al., (2016) point out that “to be conscious of more is not to be more conscious” (p. 407). An individual subject may be conscious of more or less content, and a subject may be more or less *physiologically* conscious, but it does not seem that a given adult subject can be more or less *phenomenally* conscious from time to time (*pace* Lycan (1996) and Dennett (1995), although I return to this topic in 8.6.3).

I leave open the possibility that across species or even across individuals of the same species, phenomenality may admit of degrees, or be continuous rather than discontinuous, to use Velmans’ (2012) term. Organisms other than human—say dogs or early hominids—may possibly be less phenomenally conscious (Bering & Bjorklund, 2007; Corballis, 2007; Polger, 2017). Advanced aliens or post-human species may be more phenomenally conscious. It may even be that individual

humans are comparatively more or less phenomenally conscious compared to each other, but this seems something we can never really know—it is a variation on the Problem of Other Minds. An infant may be less phenomenally conscious than an adult (Trevarthen & Reddy, 2017; Zelazo, Gao, & Todd, 2007), but it seems impossible in our current state of knowledge to determine this with any certainty.

Another interesting case is the possibility of “hyperconsciousness” due to drug effects or meditation, sometimes described in terms that might suggest increased phenomenality compared to the subject’s experience before and after the special state. But there is some reason to think this is an increase in the clarity or vividness of the experience, which is strictly a difference in the character of the content rather than the phenomenality as such (Droege, 2009, p. 83 fn. 18). Later, I consider certain aspects of clarity and vividness (3.3.3.3), meditative (6.3.3.5) and drug-induced states (6.3.6.1), and the possibility of a phenomenal state without any content whatsoever (6.3.3).

2.6 Chapter Summary

While the concept of consciousness has a long and promiscuous history, I am here most interested in *phenomenal consciousness*, characterised and identified by four inherent features: what it is like-ness; a situated first-person perspective; phenomenal unity; and temporality. A further distinction between the content of consciousness and phenomenality as such will prove important to the arguments that follow. Having established what I mean by “consciousness” in this treatise, I turn now to “attention.”

3 The Many Faces of Attention

3.1 A Brief History of Attention

The concept of attention has just as long a history as that of consciousness. It is recognisable in the ancient Stoic concept of *προσοχή* (*prosochē*)—the continuous state of being attentive to one’s present self, to the sensations, thoughts, and actions that mould and shape one’s moral character, and is essential to making progress on the path to *eudaimonia* (the good life), and ultimately, to becoming a Stoic sage (Fisher, n.d.). In the East, we find that the ancient Buddhist concept of “*manasikāra* lies at the origin of all experienced phenomena; since phenomena arise with the rising of attention” (Analayo, 2012, p. 196), highlighting the intimate connection between attention and consciousness that persuaded many thinkers through history till the present to hold the two inseparable. Buddhist concepts of attention may have been the first recorded instance of the concept being characterised using the idea of bottleneck in cognition (Adamson & Ganeri, 2017), a characterisation that played a central role in the twentieth century. The concepts of mindfulness, the focusing of one’s attention upon the experience of the present moment, and *nepsis* (*νῆψις*),

which connotes both a kind of vigilance over one's senses, thoughts, and actions are also found in ancient Christian mystical traditions (Morelli, n.d.). Interest in the concept continued through medieval and Early Modern periods to the present.³⁰

If the concept of consciousness has been difficult to define, so has that of attention, albeit for different reasons. Whereas consciousness is ontologically ineffable (or at least, elusive), attention is definitionally promiscuous—the term has been used in vast variety of ways, as we shall shortly see. This promiscuity threatens any attempt to answer Q, since any answer will depend on identifying the presence or absence of attention. How can we do this if we don't know what characterises attention? In this section, I hope to bring sufficient order to the chaos to make it possible to identify attention in interesting and useful ways, thus allowing for a useful discussion of Q in Part II. While I mostly focus on visual attention, much of what is said can, I think, safely be generalised to other modalities and cognitive functions.

3.1.1 Chapter Outline

In this chapter, I address some important preliminary issues, briefly surveying some of the definitional approaches others have taken, describing some useful distinctions between kinds of attention, and highlighting some of the difficulties one faces in finding a definition of attention. I then use examples from the literature on attention to develop my own descriptive taxonomy of definitions of attention, categories of which cut across the distinctions earlier described. Finally, I describe a particular approach to defining attention, upon which the arguments in Part II are based.

³⁰ For the history of attention, see Styles (1997, Chapter 2) and Mole (2013).

3.2 Preliminaries to Defining Attention

“No-one knows what attention is.”

~what William James *should* have written.³¹

I take it as obvious that if the terms and concepts under investigation in any field of research are not common to all researchers, it will be very easy for them to be at cross-purposes in ways that hinder the progress of that project. This is particularly evident in the study of the relationship between attention and consciousness, and is a concern expressed by many of its leading exponents. Here are a few recent examples:

“it is necessary here to define the two key concepts (attention and consciousness), as these terms have been used in so many different ways in the literature” (Iwasaki, 1993, p. 213).

“Although often used in everyday speech and in the scholarly literature, ‘selective attention’ and ‘consciousness’ lack clear definitions. Partly because of this deficit there exists a lively debate on the relationship between the two” (van Boxtel, Tsuchiya, & Koch, 2010, p. 1).

³¹ What he actually wrote, and what most authors quote, is “Everyone knows what attention is” (James, 1890), reflecting Munsell’s (1873, p. 11) less pithy but earlier statement: “On attention itself, it is needless to discourse at length; its nature and conditions are familiar to every thoughtful student.”

“precise operational definitions are likely to be necessary in order to understand the relationship between attention and consciousness” (De Brigard, 2010, p. 200).

“One thing that is clear is that progress in this debate seems unlikely if we do not pay more attention to the definitions of key terms involved” (J. H. Taylor, 2013a, p. 192).

William James’ confidence that we all know what attention is would seem misplaced, at least for cognitive scientists and philosophers. “Although we all know what it feels like to pay attention, the concept is notoriously difficult to define” (Rosenberg, Finn, Scheinost, Constable, & Chun, 2017, p. 290). Folk psychological intuitions may be all very well in general discourse, but when we come to address a question like Q, it quickly becomes evident that we need well delineated, workable definitions of the concepts under investigation.

3.2.1 Other Definitional Approaches

Attempts to solve this definitional problem have taken various paths. Few today think of attention as a discrete system with discrete neural realisers, as is, say, the visual system. Any satisfying definition of attention will have to encompass its pervasiveness across all modalities of perception, thought, and action.³²

Adverbial approaches can be traced back at least as far as Locke (Mole, 2013, sec. 1.3), who counted attention among other “modes of thinking”, such as reverie,

³² Watzl (2011b) discusses a number of reductive and anti-reductive approaches to capturing the nature of attention.

intention, and study. Indeed, many approaches to attention see it as *metacognitive* (Baars, 1988, pp. 302–303)—a way of thinking, rather than thinking itself. Mole’s *Cognitive Unison* model of attention is explicitly adverbial: “What is essential to an adverbial phenomenon is not what happens to what but how the things that happen happen” (Mole, 2011b, p. 70).

Defining attention as an adverb at best tells us only the broad ontological category to which it belongs, but it does not yet tell us how to identify attention in the field. That will require a lot more work. Two broad approaches commend themselves. De Brigard & Prinz (2010) contrast the *commonsense* approach to the *stipulative* approach.

The *commonsense* approach is based on the idea that we just know attention when we see it. William James famously wrote, “Everyone knows what attention is” back in the nineteenth century, although it has been suggested with good reason that he might have had his tongue in his cheek when he wrote it (Hardcastle, 1997; Mole, 2011b, p. 60). In the twenty-first century, “The concept [of attention] is still introduced at present-day conferences with introspective examples from everyday life that show attention at work when looking for a friend’s face in a crowd, or when searching for a particular pen on a cluttered desk” (Ruff, 2011, p. 1). For Prinz (2011, p. 183), attention is a “natural kind term ... not something that has an essence that can be discovered by conceptual analysis”. He defines attention by examples of when it takes place: search and pop-out; monitoring; tracking, vigilance; selection; and survey. He goes on to say, “I don’t think any of these phenomena constitute a definition of attention. Rather they are all cases in which we say that attention is taking place.” Thus, we somehow just “know” attention when we see it, and similarly know how to recognise its absence.

But as De Brigard and Prinz (2010, p. 51) rightly point out, one of the problems with a commonsense approach is that it is very hard to pin down. Folk psychology is fraught with diversity and imprecision. A recent debate concerns sleeping mothers who wake up when they hear their babies cry. Does this constitute attention leading to consciousness, or consciousness leading to attention? This is another example of

how two perfectly reasonable yet contradictory commonsense answers can be supported (De Brigard, 2010; De Brigard & Prinz, 2010; Mole, 2008a). What is more, De Brigard's (2010, pp. 195–199) analysis of how one might judge a particular cognitive process as being attention or not highlights the influence other situational factors can have on one's conclusions, much as our judgements of whether or not the colour of a dress is blue depends heavily on how it contrasts with other nearby colours to which we might compare it.

On the other hand, a *stipulative* approach is one where we merely create artificial borders that constrain which cognitive processes shall be stipulated to be denoted by the term, "attention." But this approach too has its problems. Not everyone agrees where the joints lie. Creating artificial borders runs the risk of begging the question of the relationship between attention and consciousness: if one begins by stipulating that attention is directed all and only to that of which a subject is conscious, then it is unsurprising (and uninteresting) that one can then conclude that attention is both necessary and sufficient for consciousness. As De Brigard and Prinz point out, one ought to beware of "idiosyncratic definitions that settle crucial questions by fiat rather than facilitating the process of scientific investigation and discovery" (p. 52).³³

Another approach is to embrace the diversity of ways the term attention has been used and accept that "The vast literature on attention makes it clear that it is not a unitary process" (Mack & Rock, 1998, p. 25).³⁴ Attention can be taken to be an *umbrella* term: an overarching concept that encompasses a broad variety of other concepts. This is an idea that has been around for some time:

³³ Barrett (2014) and Taylor (2013a) argue that De Brigard and Prinz themselves ultimately fall into this trap, despite expressing an awareness of the danger.

³⁴ See also Allport: "Even a brief survey of the heterogeneity and functional separability of different components of spatial and nonspatial attentional control prompts the conclusion that, qua causal mechanism, *there can be no such thing as attention*. There is no one uniform computational function, or mental operation (in general no one causal mechanism) to which all so-called attentional phenomena can be attributed" (Allport, 1993, p. 203, italics in the original).

“Nor again, I should add, is there any one special activity at all, but various activities, if they lead to one result, are called attending” (Bradley, 1886, p. 305).

But the nature of this umbrella is not so clear. How are the different uses of the term attention related to one another? Are they something like the elephant of fable?

“the field and literature lack a common language to communicate and connect their work with each other. Drawing analogy with another folk story, to abandon the term attention would cause all blind men or women to feel different parts of the elephant, not realizing that they are touching the same animal” (Chun, Golomb, & Turk-Browne, 2011, p. 91).

It would therefore be helpful to have a principled way to classify all the different items under the umbrella and—more appositely to Q—to distinguish between attention and the absence of attention.

3.2.2 My Approach

While a definition ought ideally to capture the essence of the thing it defines, ultimately, it is a tool intended for a certain purpose. Different approaches to defining a concept may better suit different purposes. Purpose is not the sole arbiter of the validity of a definition, of course. Stipulation can be taken too far. But nature rarely imposes extremely narrow restrictions on how it allows itself to be carved up.

“Scientific terms are not given in nature. They are worked out, often over generations of data collection and debate” (Baars, 1997b, p. 363). In this section I take neither a commonsense nor a stipulative approach, but a *descriptive* approach, identifying and categorising some of the most interesting and common phenomena that have been allocated under the umbrella of “attention” in the literature.³⁵ From this, I extract what it is they have in common to develop an approach to the concept of attention that I hope will be useful not only in this treatise, but perhaps for other researchers as well.

This descriptive approach does not require me to endorse any of the definitions as definitive. It is closely related to the approach biologists take to grouping and taxonomizing living creatures. Like biological taxonomies, my attentional taxonomy is not set in stone, but liable to improvement; and it is not the only way to carve up attention, although I believe it is the most helpful way for addressing questions like Q. But first, some important distinctions and difficulties.

3.2.3 Some Distinctions

A number of distinctions have been drawn among different types or varieties of attention. These distinctions cut across the categories I develop below in my taxonomy of attention (3.3), but it will be useful to describe them briefly here first (Table 1). A more detailed account of these distinctions may be found in Appendix 1,

³⁵ I take my inspiration from Samuel Johnson (1709-1784), whose pioneering work captured the contemporaneous use of language in eighteenth century England in the remarkably influential *A Dictionary of the English Language* (1755). His method consisted chiefly of gathering examples of the many different usages of a word from existing literature on slips of paper and collating them. I cannot claim to match either his patience or the breadth of his reading (and certainly not his wit), but then again, my subject—attention—is just one entry in a dictionary.

including some discussion of some of the more controversial aspects of these distinctions.³⁶

These distinctions arise regularly in the literature on attention and Q, and a clear idea of them will prove helpful in navigating the arguments in Part II. But the distinctions do not constitute a definition or a principled way of identifying the presence or absence of attention. Before I attempt to provide just that, I briefly consider some of the difficulties that have hitherto afflicted others who have also made the attempt.

Table 1. A summary of some important distinctions among kinds of attention.

<i>Where is attention focused?</i>	Internal Attention Attention to internal states	External Attention Attention to the external world
<i>How is attention recruited?</i>	Endogenous / Top-Down Attention Attention is directed by the subject	Exogenous / Bottom-Up Attention Attention is drawn by external stimuli
<i>Is volition involved?</i>	Voluntary Attention I choose to attend	Involuntary Attention My attention is drawn without my choice
<i>Are there degrees of attention?</i>	Focal Attention The “centre of attention”	Background Attention Partial attention to the periphery of what is attended

³⁶ For a briefer and slightly different list of distinctions, see Watzl (2011b, p. 846).

<i>What is being attended?</i>	Spatial Attention Attention to a region or locus in (for example) a visual field	Object Attention Attention to a composite object as a whole Feature Attention Attention to specific features of an object, rather than the object as a whole
<i>How is attention related to time?</i>	Synchronic Attention Attention at a moment in time	Diachronic Attention Attention over or through time
<i>At what level is the attention?</i>	Personal-Level Attention The subject attends	Sub-Personal-Level Attention Some subsystem or cognitive process attends

3.2.4 Difficulties Defining Attention

Defining a term that has been employed as promiscuously as “attention” has is a difficult task (Mole, 2011b). Three issues in particular are relevant to Q, so below I briefly describe them, in the hope that I will adequately deal with them in my own treatment of the question. The three issues are: *“merely verbal disputes”*; *question begging*; and *explanatory burden*. One way to address these issues is to develop a *taxonomy of attention*, and I conclude this section with a review of the few existing taxonomies, before developing my own in the next section.

3.2.4.1 “Merely Verbal Disputes”

“if there really is to be a definitive answer to the question of whether attention is necessary and sufficient for consciousness, then obviously we are not free to define ‘attention’ in any way that we choose. If we simply claim that the theorists in question are just working with different concepts of attention, and that this is simply a linguistic issue, then we are dangerously close to saying that ‘really’ there is no answer to the question of whether attention is necessary and sufficient for consciousness, because different theorists will deliver different answers depending upon their definition of ‘attention’” (J. H. Taylor, 2013a, pp. 190–191).

To avoid the trap Taylor highlights,³⁷ we need an effective way to avoid descending into merely verbal disputes (Chalmers, 2011; Seager, 2016, p. 207).³⁸ My solution is to develop a descriptive taxonomy of how attention is employed in the literature, acknowledging the validity of these different senses of attention for different purposes and in different contexts, then distilling the concepts that tie them together to get at the essence of attention. This can be seen as an application of Chalmers’ (2011, p. 526ff) *method of elimination* in that I seek out what exactly people are describing when they use the term “attention,” and base my arguments on *that*, diverse as it is. Out of this approach emerges a concept attention that is expedient for addressing Q in interesting and useful ways.

³⁷ See also Kentridge (2011, p. 229) and Smithies (2011, p. 247).

³⁸ A parallel problem arises in phenomenology (Kriegel, 2007, pp. 123–124).

3.2.4.2 *Question Begging*

“The trouble is that if one starts off with a definition that determines what the evidence is saying, finding out what the relation between consciousness and attention is will no longer be an empirical puzzle. And if one thinks, as we do, that this is an empirical problem, one has to be very careful in defining the terms in such a way that neither of them implies analytically the other” (De Brigard & Prinz, 2010, p. 52).

There is a very real danger in exploring Q in defining attention *a priori* in such a way that it is inseparable from consciousness, and thereby begging the question of whether they can come apart, and rendering moot any empirical search for one without the other.³⁹ This is a failing of Phenomenal definitions of attention (3.3.3) that has historically hampered progress on Q (or its antecedents). I take it as obvious that attention must be assumed to be conceptually independent of consciousness, and the relationship between them to be a matter for empirical investigation (1.4.1).

3.2.4.3 *Explanatory Burden*

If we manage to avoid both merely verbal disputes and question begging, we are still left with the problem of making sense of all those different approaches to defining attention. Mole (2011b, p. 61) argues that the difficulties involved in defining attention may derive from the fact that we are asking a single term to cover far too

³⁹ See also Barrett (2014, p. 9,10), Mole (2008a, p. 93).

many things, much like Bilbo Baggins' sad predicament: "I feel thin, sort of stretched, like butter scraped over too much bread."⁴⁰

Mole discusses three possible solutions to this *explanatory overburdening* problem.⁴¹ First, we might give up altogether on the idea of a single, unitary process called "attention" and instead adopt a *piecemeal* solution, breaking down attention into many related yet discrete processes. But this fails to respect the commonality that binds together all those things we call attention. Second, we might adopt an *eliminativist* solution, and simply deny that there is such a thing as attention. The folk psychological view is simply mistaken, and there is no causal mechanism to which the term "attention" refers, and therefore it is fruitless to try to identify, much less understand it. But this too is unsatisfactory. Not only does "attention" play a valuable role in daily life, it is both a highly useful construct for empirical research (and has been the subject of much research), and a very plausible, indeed, integral part of good psychological models of cognition. Recognising that it is not a simple unitary thing, but a complex mishmash of sub-things in no way diminishes the usefulness of the umbrella term "attention." Third, one might take a *something else* approach—if attention is not a causal mechanism in the brain, then perhaps it is something else. Mole's own adverbial approach might be an example of this solution.

3.2.4.4 *The Need for a Taxonomy of Attention*

My descriptive approach below builds upon some of the lessons from Mole. To identify the presence or absence of attention, we need *not* know what attention *is*. It is enough to—as with consciousness—be able to recognise its *features*. But the diversity of ways "attention" has been used makes this task fraught. The list of

⁴⁰ J.R.R. Tolkien, *The Fellowship of the Ring*. See also Watzl (2011b, sec. 5).

⁴¹ See also Fazekas and Nanay's (2018) discussion of and solution to the "disunity problem," Henry Taylor's (2018) argument for favouring pluralism over monism, and Ganeri's (2017, p. 222) argument for favouring a "varieties of attention" account over "attention essentialism."

features that identify attention is going to be much longer, and far more complex than that for consciousness. Yet such a *taxonomy of attention* would be invaluable in resolving the explanatory burden problem that arises from definitional diversity and dissolving the apparent disconnection between the monistic accounts of attention of the philosophers and pluralistic accounts of attention of the psychologists (H. Taylor, 2018). Does any such taxonomy exist?

“Although the diversity of attention is recognized, it is also true that no completely satisfactory taxonomy of attention has been put forward” (Parasuraman, 1998, p. 5).

“At first, attention was assumed to be a unitary phenomenon. More recently researchers have increasingly commented on the absence of a convincing taxonomy of attention (or even a useful definition) and the heterogeneity of experimental findings, suggesting several different attentional mechanisms” (Rees, 1999).

There are not many attempts at the explicit taxonomisation of attention in the literature. Treisman (1969) describes four kinds of attentional strategy: restriction of the number of inputs; restriction of dimensions analysed; looking or listening for specific items; and selection of outputs for action and storage in memory. My taxonomy below is a significant expansion on Treisman's. A notable recent taxonomy is that of Rensink (2015), who adopts a more disciplined approach to the question of taxonomising attention, and proposes some guidelines for the project (pp. 349-350). In brief, his taxonomy uses function, perceptual effects, and mechanism to categorise varieties of attention. On this basis, he proposes a simple taxonomy, involving an initial division of attention into two functions, with each of these subsuming some sub-functions: *orientation* subsuming *sampling* and *filtering*; and *integration* subsuming *binding*, *holding*, and *individuating*. My taxonomy leaves out binding and holding for reasons I discuss below (3.4.2), but includes a much broader range of functions and some non-functional approaches.

Rensink's taxonomy has a number of valuable features. First, it connects sub-functions of attention to their experimental or empirical perceptual effects. For example, attentional blink and change blindness can best be explained as features of how attentional holding works (p. 357). Second, it attempts to make some connections ("dependencies") between the sub-functions of attention (p. 360), a complex project I will not attempt here with my broader taxonomy. Rensink's taxonomy of visual attention is one of the few in the literature. Others are more specialised, such as Regan et al's (2011) taxonomy of Driver Distraction and Inattention, which cuts across many of the categories in Rensink's taxonomy. Arvidson (2003) provides a "lexicon" that translates cognitive science concepts related to attention into phenomenological language, thereby covering some of the territory I cover below, and elsewhere in this thesis.

None of these taxonomies is comprehensive enough to provide a solid basis for the identification of attention relevant to Q, and indeed, none was designed specifically for that purpose. Well, we shall have to invent one.

3.3 A Taxonomy of Attention

"It is certainly premature to claim that neuroscientists have understood all properties of neural, cognitive and sensory processing that define an instance of attention. But it appears that decades of intensive research have resulted in a body of work that may allow us to formally define attention from a mechanistic neural perspective. Such developments are encouraging, and may ultimately help us to understand which subdefinitions of the

concept of attention may be justified by biological reality. Time will tell whether such notions will venture outside of neuroscience, to complement classic introspective definitions of attention based on verbal descriptions and mental states” (Ruff, 2011, p. 18).

In this section I survey the literature to build a detailed taxonomy of definitions of attention, which serves as the basis for my own approach to identifying attention in a manner capable of providing interesting and useful answers to Q.

As Ruff observes, attention is not a single, monolithic mechanism or process. Any account of attention will have to do justice to *all* its “subdefinitions,” as an incomplete taxonomy of attention casts doubt on the identification of the presence or absence of attention. My approach, then, is to seek “unity in diversity.” I embrace the rich variety of extant definitions of attention and then distil the essence of what ties them together (3.4.1). These definitions are not limited to any one sensory modality, or even to perception itself, but may apply across all cognition.

Although I attempt a comprehensive taxonomy, I make no guarantees of succeeding perfectly in said attempt. The literature on attention is huge, and I have not been able to exhaust it.⁴² Further, many of the entries overlap conceptually, and some of them could plausibly be categorised under a different heading. For these reasons, my taxonomy is liable to revision and improvement. Nor is my taxonomic structure the only one possible.⁴³

⁴² See also the brief “list” of ways of defining attention in Mole (2011b, pp. 62–63) and the taxonomy of Chun et al., (2011).

⁴³ Another taxonomic distinction that might be made is to divide definitions of attention according to whether they are at the level of the whole person, or of particular cognitive processes, or at the neural level. Some definitions describe what is thought to constitute attention, whereas others describe operational features. Yet another distinction might delineate which of the senses is involved. Or attentive processes might be categorised according to the level of processing at which they occur—e.g., at the level of individual visual characteristics like edges, or at the level of bound visual

I present the Taxonomy in table form first (3.3.1, Table 2) to give the reader an overview, and then proceed to assemble it from the literature. To help illustrate each entry, I use the example of the *invisible gorilla*. In the famous *inattention blindness* experiment of Simons and Chabris (1999), subjects watch a video of basketball players passing the ball to each other, and are asked to count the passes. Many subjects do not notice that a man in a gorilla suit walked right through the middle of the scene while they were counting.

3.3.1 A Definitional Taxonomy of Attention

Table 2 below summarises the taxonomy of definitional approaches to attention that I have developed above. However, it should be noted that while a table is the simplest format, it is less than ideal, for there is a great deal of similarity, connection, and overlap between many of the approaches.

Table 2. A definitional taxonomy of attention.

- 1. Behaviourist**
- 2. Phenomenal**
 - A. Chief Tenancy of Consciousness
 - B. Confidence
 - C. Clarity and Vividness
 - D. Saliency

objects, etc. All these other taxonomic distinctions cut across the particular categories I have chosen here.

- E. Fixedness
- F. Directedness

3. Mechanistic

A. Parts

- I. Anatomical Regions
- II. Connectomes

B. Operations

I. Access

- a. Simple Access (but not accessibility)*
- b. Abundance /*
- c. Increase / Amplification*
- d. Maintenance*
- e. Influence / effects*

II. Detection and Recognition

- a. Orienting*
- b. Vigilance or Monitoring*
- c. Scanning or Searching*
- d. Expectation*
- e. Alerting*
- f. Tracking*

III. Selection

- a. Filtering*
- b. Gateway*
- c. Spotlight*
- d. Zoom In / Out*
- e. Exclusion / Suppression / Absorption / Dedication /
Absence of Distraction.*
- f. Competition*
- g. Selection for Action*

IV. Control

V. Coherence

C. Organisation

- I. Cyclic pathways

- II. Feedback loops
- III. Contextual interconnection

3.3.2 Behaviourist Definitions

Behaviourism is, roughly, the positivist view that we can really only know about observed behaviours, but we can have no useful access to what is going on inside the head to produce those behaviours. Today, philosophical Behaviourism has fallen somewhat out of favour, but there have been authors who formulated a definition of attention in Behaviourist terms. For example, in the nineteenth century (before the advent of the Behaviourist movement) Theodule Ribot considered the outward marks of attention to be “the necessary conditions, the constituent elements, the indispensable factors of attention” (Ribot, 1890, p. 19).⁴⁴

This intuition is a strong one and there is something to be said for the importance of the consequences of attention as being part of that which defines it. In practice, that is the most common way we recognise attention in others, and occasionally, even in ourselves. Take the invisible gorilla example. We identify the absence of attention in the test subjects by the fact that they are unable to verbally report the presence of the gorilla, and perhaps by observing that their eyes track the passes but never turn to the gorilla—all of which are behavioural markers.

So, a distinction between *behavioural* markers or definitional characteristics of attention, and *Behaviourist* definitions of attention is one worth making. While the excesses of the Behaviourist approach seem unnecessary, for all the reasons that Behaviourism has lost favour, the behavioural approach continues to play an

⁴⁴ See also a discussion of his view, and comparison with those of Alexander Bain, in Mole (2013, sec. 1.5).

important role in attention research, as when, for example, attention is defined as occurring just when a subject meets certain experimental criteria, or passes a certain threshold on an experimental scale of attention. But we shall see in 7.4 that this approach has some serious limitations. Thus, while formulating Q in Behaviourist terms is unfruitful, respecting the behavioural dimension of Q is inevitable and useful, so long as due care is taken.

3.3.3 Phenomenal Definitions

Behaviourism leaves out something important. Often when we attend, there is something it is like to attend. But is attention always phenomenal? On some definitions, it is. What I am calling *Phenomenal* definitions of attention are those approaches that include a phenomenal aspect of attention as a *sine qua non* feature of attention: there is *always* something that it is like to attend. In terms of phenomenal experience, Phenomenal Attention is the antithesis of Behaviourist Attention—it grounds attention in the phenomenality that Behaviourism dismisses.

The chief problem with Phenomenal Attention, of course, is that it defines the answer to Q *a priori*, rather than allowing Q to be an open empirical question. If consciousness lies at the heart of attention, then it is necessary for attention. Whether it is sufficient will depend on what one thinks of the conscious content outside the focus of attention—whether it is unattended, or merely poorly attended. I have grouped Phenomenal definitions in six sub-categories: *chief tenancy of consciousness; confidence; clarity and vividness; salience; fixedness; and directedness*.⁴⁵

⁴⁵ For a summary of some of these subcategories of Phenomenal Attention, including citations to Early Modern and Modern examples, see Watzl (2011b, p. 843).

3.3.3.1 “Chief Tenancy of Consciousness”

That which is attended is just that of which we are most conscious. Thus, the subjects cannot report the presence of the gorilla because its image never dominated their phenomenal content, and they can only report that which does so dominate.

“Take, for instance, B.F. Bradley’s old view. He starts off saying that people have taken attention to be ‘predominance in consciousness’, and he cites J.S. Mill in support: ‘The expression [attention] means that a sensation tends more or less strongly to exclude from consciousness all other sensations’. Thus, accordingly, Bradley decides to call attention ‘a state which implies domination or chief tenancy of consciousness’ (Bradley, 1886, p. 306).⁴⁶ In light of this definition, then, attention involves the activity of consciousness; indeed, it implies that there cannot be attention without consciousness, i.e., that consciousness is necessary for attention” (De Brigard & Prinz, 2010, p. 52).

The prominence of some object of attention over all other possible objects of attention within the subject’s cognitive economy, and the exclusion of other possible objects, are features that belong to non-phenomenal definitions below (3.3.4.2.3). However, the defining characteristic of this approach is that it emphasises the phenomenal aspect of these intrinsically non-phenomenal activities—there is *something it is like* to perform them. This approach bears some interesting similarities to contemporary accounts that define attention as physiological consciousness or “maintaining an alert state” (e.g., Posner, 1994, p. 7398). The

⁴⁶ De Brigard and Prinz cite this mistakenly as “(1886, 22)”. I have corrected the page number.

interesting feature of this definitional approach for our present purposes is the way it describes attention as a *property of consciousness*, thus determining the relationship between the two by an act of definition.

3.3.3.2 *Confidence*

The definition of attention as *confidence* can be traced back at least as far as Descartes:

“So long as we attend to a truth which we perceive very clearly, we cannot doubt it. But when, as often happens, we are not attending to any truth in this way, then even though we remember that we have previously perceived many things clearly, nevertheless there will be nothing which we may not justly doubt so long as we do not know that whatever we clearly perceive is true” (Descartes, 1988, p. 309).

Whatever we experience without attending is relatively vague and unsure.⁴⁷ Paying attention to a thing brings a greater degree of certainty to the experience of that thing, and this is reflected as the phenomenal sense of confidence. Mole concludes that for Descartes, “the move from radical doubt to certainty about the truth of particular clear and distinct ideas—is, therefore, a transition that is mediated by attention” (Mole, 2013, sec. 1.1).

The subjects in the invisible gorilla experiment did not report the gorilla because they failed to attend to it in a way that would have given them a confidence about its

⁴⁷ This point becomes very important in Chapters 7 and 8.

passing through the scene. Interestingly, Simon & Chabris (1999) tried to tease out a report of the gorilla from subjects using not one question, but a series of four, gradually more explicit questions, starting with “did you notice anything unusual?” and ending with “did you see a gorilla walk across the screen?” (p. 1068). Yet only one out of nearly two hundred subjects showed any kind of wavering through this series of questions. That is, nearly all the subjects were quite confident in their answers, whether yes or no, regardless of the suggestive line of questioning. On this definition of attention, this confidence tracks perfectly where the subjects’ attention was directed, whether to the gorilla, or not. On the other hand, Matthews et al., (2018, p. 5) found that subjective confidence ratings track objective *accuracy* rather than attention.

3.3.3.3 *Clarity and Vividness*

While confidence is about a propositional certainty *that* something is thus and not otherwise, there is a closely related Phenomenal definition of attention that defines it as a perceptual *clarity* or *vividness*.⁴⁸ In practice, it is hard to imagine vividly seeing a bright red rose and not being confident that you see it, but the two can come apart at least conceptually. Without doubt, attention usually enhances clarity and vividness. The vague and shadowy figure I might have felt flitting around my visual field becomes a clear and vivid (if somewhat unconvincing) gorilla the moment I turn my attention upon it. Even if I look upon a visual scene through the out-of-focus lens of a camera, I can experience the very lack of proper focus clearly and vividly if I attend to it.⁴⁹

⁴⁸ This distinction between reflective and propositional confidence on the one hand and perceptual clarity and vividness on the other plays an important role in 7.2.1.3.

⁴⁹ For an argument against the suggestion that vividness can vary independent of content, see Bourget (2017).

This idea of attention as clarity seems to have attracted interest in the early twentieth century,⁵⁰ but has fallen largely out of favour in modern times. Treisman (1964, p. 12) laments that the idea of attention as “the increased clearness of a particular idea” has proven to be sterile in psychological research, and Watzl (2011a, pp. 151–152) argues against the very similar perceptual “determinacy view.” But the idea has continued to play a role in at least some research. For example, Baddeley and Andrade (2000) showed that selectively taxing the working memory of a sensory modality attenuated the “phenomenological vividness” of perceptions in that modality. And Schlagbauer et al., (2018) found that contextual cueing—a way to draw a subject’s attention to the configuration of visual display elements—enhances the “clarity” (p. 2) of the subjective experience of both the target object and the surrounding visual configuration.

However, there is some reason to doubt that attention alone is the sole determinant of phenomenal clarity and vividness. For example, Wassell et al., (2015) found that higher blood (and salivary) progesterone concentration correlates with an enhancement of the vividness of voluntary visual mental imagery, which is perhaps in keeping with other research that suggests that females tend to have more vivid imagery than males (Campos & Pérez, 1988). However, the mechanism of the enhancement has yet to be elucidated. The possibility remains that elevated levels of progesterone might exert their enhancing effect upon vividness via enhancing attention, as adrenaline (epinephrine) most likely does in a classic fight or flight response.

⁵⁰ Stazicker (2011a, p. 172 footnote 12) describes Leibniz’s ideas of clarity and vividness in relation to attention and consciousness. Philips (2011a, p. 221, note 4) observes that discussions of attention as “clearness” may be found in authors such as Titchener, Woodward, and Gill and Dallenbach in the early twentieth century. A modern revival of the idea is Jennings (2012).

3.3.3.4 *Salience*

Salience has been described as an “attend to me” signal (Sawaki & Luck, 2010). Potential targets of attention are more salient just insofar as they are more likely to draw attention. Salience is generally characterised as bottom-up attention—features that “stand out” are more salient. But in fact, it is likely to be a complex interplay between bottom-up standing out and top-down context-sensitive biasing (Egeth, Leonard, & Leber, 2010, p. 130).⁵¹ Whereas a gorilla in a basketball game would normally be highly salient, the top-down attention to following the basketball overpowers this bottom-up signal. As a Phenomenal definition of attention, the salience is about that phenomenal sense of urgency or importance that—whether bottom-up or top-down—draws attention to a target. In the positive symptoms of schizophrenia, this sense ectopically results from unimportant stimuli (Fletcher & Frith, 2009).

Ruff (2011, p. 5) observes that salience is readily capable of being characterised in the non-phenomenal terms of neural processing patterns: a bottom-up effect of one or more stimuli evoking a stronger neural response than other stimuli competing for limited neural response resources (compare Selection, 3.3.4.2.3). Indeed, much of the attentive processing of salience occurs subconsciously, and it is possible for unconscious objects—masked nudes, for example—to attract attention-as-salience (Jiang, Costello, Fang, Huang, & He, 2006). This intuition is further supported by work such as that of van Swinderen (2005, p. 324) which suggests that salience mechanisms might operate in producing attention in fruit-flies, whose capacity for phenomenal consciousness remains an open question (Barron & Klein, 2016; Key, Arlinghaus, & Browman, 2016; Tiffin, 2016).

⁵¹ For a discussion of the different roles played by salience and attention in this dance in the context of making moral decisions, see Chappell and Yetter-Chappell (2016, p. 454). For a discussion of the neural level at which salience operates in the mechanism of attention, see Fellrath and Ptak (2015) and Parkhurst et al., (2002), and for a predictive coding account, see Clark (2016, pp. 28, 66–69).

3.3.3.5 *Fixedness*

Fixedness is the experience of being restricted to a narrow train of content such that one is unable to escape that train and enter into other trains. The subjects' task of counting the passes fixes their attention narrowly on the ball and the players, excluding interpretations of the scene that include an ectopic gorilla. Baars (1988, pp. 143–145) develops this view of attention by considering a sentence like the following:

“The ship sailed past the harbour sank.”

Most readers will initially have trouble making sense of this sentence, until they realise it may be read as, “the ship—which was sailed past the harbour by someone—sank.” Before this realisation, the reader's interpretation of the sentence is *fixed* by the assumption that the most obvious subject of the verb “sailed” is the ship, rather than the people sailing the ship. Like top-down imperatives in salience, context is an important dimension of fixedness here for Baars. Attention-as-fixedness is having our experience trapped by powerful hierarchies of context in a particular way of experiencing to the exclusion of other possible ways of experiencing. Thus, fixedness is an inherently *phenomenal* form of attention. For example, “in absorbed states of mind—in reading an engrossing novel or watching an entrancing motion picture—we are deaf and blind to the world” (p. 145).⁵²

⁵² Jennings (2015, pp. 288–289) discusses an interesting development of this approach in which she speaks of “focus without the aid of attention”.

3.3.3.6 *Directedness*

Related to Baars' fixedness, the idea of *intentionality* (Menary, 2009) has also been conscripted to the task of defining attention. Here, one attends to a thing just when one *directs* their thoughts to that thing. For example, understanding a word might be thought intuitively to imply both attending to that word and being conscious of the word and its meaning. The subjects' attention (and therefore the content of their consciousness) is constituted by their directing their thought towards the movement of the basketball, but never to the gorilla. Once again, attention and consciousness are tied together by definition. For example, in a paper suggestively titled "Attention: The Mechanisms of Consciousness," Posner (1994, pp. 7400–7402) discusses attention as "attending to ideas." Some more recent accounts that might fit plausibly under this subcategory are Smithies' (2011) rational-access view and Koralus' (2014) erotetic (question-related) theory of attention.

Phenomenal Definitions of attention satisfy the powerful intuition that there is a close connection between attention and consciousness, but they fall into the trap of assuming *a priori* that they are inextricable from each other, a trap best avoided in the quest to answer Q (3.2.4.2). The third class of definitions avoids that trap.

3.3.4 Mechanistic Definitions

If Behaviourism defines attention solely by its outward markers, a more contemporarily popular approach is to define attention by its internal markers—patterns of brain activity. Using Bechtel's (2008, pp. 13–17) model of mental mechanisms, my *Mechanistic* definitions identify attention by the *parts* in the brain that subserve it, the *operations* those parts perform, or the ways that those parts and operations are *organised*.

3.3.4.1 *Parts*

Attention might be defined as activity in the brain structures that subserve it. But there are at least two ways to delineate these parts: *anatomical regions* and *connectomes*.⁵³

3.3.4.1.1 *Anatomical Regions*

There has been some progress made on identifying brain regions that are involved in specific types of attention. For example, the fronto-parietal regions have been implicated on fMRI studies with top-down visual, auditory, and tactile attention signals, while bottom-up attention seems to be subserved by localised modality-specific regions, such as V1 for visual saliency (Kanwisher & Wojciulik, 2000). But few have been tempted to use anatomical localisation as an identifier of attention—e.g., “if there is no unusual fronto-parietal activity, the subject is not attending.” Attention is not localised in the brain: “Scientific research suggests that the class of such subpersonal attentional processes is large, highly diverse, and not well localized in the brain” (Watzl, 2011a, p. 163).

There is good reason to think that the vast majority of cognitive processes are related to brain areas in a many:many relation (M. L. Anderson, 2010; M. L. Anderson, Kinnison, & Pessoa, 2013)—each function requiring many areas, and each area subserving many functions. Attention involves many brain regions and overlaps with the footprint of many other processes (M. L. Anderson et al., 2013; Naghavi & Nyberg, 2005; Rosenberg et al., 2017). What is more, very similar acts of attention, such as task switching, may involve very different parts of the brain

⁵³ Much less likely ways are by type of neurone, or by predominant neurotransmitter (e.g., Schmitz & Duncan, 2018).

(Wager, Jonides, Smith, & Nichols, 2005). Given all this, an anatomically circumscribed localised “attention centre” is highly implausible.⁵⁴

3.3.4.1.2 *Connectomes*

An alternative to this anatomical region approach might be a network approach that carves up the brain according to *connectomes* of neurones that communicate heavily with each other, even though they are distributed throughout many anatomical regions of the brain. Mogensen and Overgaard’s (2018) *Reorganization of Elementary Functions* (REF) framework is a promising account that respects the many:many relationship and accounts for both functional localisation and the apparently contradictory capacity of brains for recovering abilities after trauma by conscripting completely different anatomical and neural structures to perform the lost functions. Thus, there are both anatomically localised functional units—“elementary functions”—and long-range connections between these units. Such an account is more felicitous given that attention of some kind likely permeates most cognitive processes (3.4.4).

Anatomical accounts of attention are relatively easy to probe with lesion studies, since lesions in the brain tend to be anatomically localised. But connectomes snaking through large swathes of the brain are virtually impossible to lesion in isolation from their surrounding neurones, even if we could identify them in a subject *in vivo*. Nor are they lesioned discretely by disease or injury. And while we have imaging tools of reasonable sensitivity to explore the activity of anatomical regions—fMRI for example—we lack satisfactory tools for selectively measuring connectome activity *in vivo*. EEG (Eimer, 2015) and MEG (Baillet, 2017) can give some

⁵⁴ Naghavi and Nyberg (2005) review a number of empirical studies that found correlations between specific brain regions and specific kinds of attention. They also found that there are areas of the fronto-parietal regions that seem increase their activity during all four of attention, conscious visual perception, working memory, and episodic memory retrieval. The significance of such findings continues to be debated.

information, but it is of frustratingly low resolution. What is more, there remain serious problems with interpreting scans (Carp, 2012) and explanatorily bridging neural activity with cognitive functions generally (Fine, 2010, pp. 281–282; Naselaris et al., 2018, p. 3).

So, defining attention by the parts of the brain involved seems unlikely to take us very far, not only because attention is so widespread in brain activity, but also because of our investigative limitations. If the parts approach is as yet unsatisfying, considering brain operations is more promising.

3.3.4.2 *Operations*

“In cognitive science, attention is usually defined in terms of its functional role, rather than its phenomenology” (Smithies, 2011, p. 250).

The *operations* of the brain can profitably be viewed as functional roles.⁵⁵ The challenge here is to distil those functional roles that are specifically best considered as attention, rather than something else (on which, see 3.4.2).

3.3.4.2.1 *Access*

Attention may be defined as simple *access* to data or content, or variations thereof (*abundance, increase, amplification, maintenance, influence, effects*), but not, I argue, mere *accessibility*. On the view that attention is just simple *access*, whatever is being

⁵⁵ For a similar, though less detailed or structured list, see Ganeri (2017, p. 224).

processed is therefore being attended. This approach suffers the serious drawback of making attention merely synonymous with cognition, and therefore a superfluous term, so it is not surprising that it finds few proponents in modern discourse. The subjects undeniably have some sort of access to the gorilla—its image falls upon the retina and must therefore register in the earliest stages of visual processing—yet they clearly do not attend to it at all.

A more popular approach is that which employs the concept of a relative *abundance* or *increase* in some cognitive quantity. Kanai et al., (2006, pp. 2334–2335) assume that an effect of spatial attention is to increase the activity generally of the orientation-perceiving circuits relating to that location in a visual field. *Abundance* is the idea that a process that recruits more cognitive processing power is attended while one that recruits less is not. Perhaps there may be a threshold,⁵⁶ above which a process is deemed to be attended, or it may be a graded affair of more or less attention.

Increase is the idea that attention is just the act of a process recruiting more cognitive processing resources. Thus, the gorilla is not attended unless its retinal image in the early stages of visual processing is passed on to later stages, or passes beyond a certain stage of processing, or exceeds some threshold of processing resource use. *Amplification* of input signals (Fazekas & Nanay, 2018) is a neural version of increase.

The thing that is in abundance or that is being increased need not be just ‘processing’ generally, but may be defined more narrowly as a certain kind of processing. For example, attention may be the increase of “access to conscious experience” (Baars, 1988, p. 302), or to particular contents of conscious experience. Or the increased access may be to unconscious data, since the question of the relationship to phenomenal consciousness is left open. Attention-as-access makes

⁵⁶ See Dehaene et al., (2006) for a similar threshold model relating to consciousness rather than attention.

attention virtually synonymous with Block's (1995) "access consciousness", which can also be separated from "phenomenal consciousness" (7.2.2).

One way to facilitate increased access or processing is the *maintenance* of access over periods of time (Kane & Engle, 2002). Maintenance in the form of providing access to content by storing it in a short-term buffer plays an important role in working memory function (5.4.1), and is one of the putative mechanisms of the capacity limitations of working memory (8.2.4).

Of course, maintained or increased access or cognitive processing implies that the outputs of attended processes will also be more *influential* or have a greater *effect* on the cognitive economy. This idea is at the heart of biasing accounts of attention (Desimone & Duncan, 1995), and it has been suggested that attention influences the character of perceptions (Carrasco & Barbot, 2019; Carrasco, Ling, & Read, 2004; Ling, 2012) and interactively influence motor movements (Moore, Armstrong, & Fallah, 2003).⁵⁷

There is a distinction to be drawn between actual access and merely potential *accessibility* (*availability*). Some have thought that attention may be defined in terms of either, or both (Chalmers, 1997; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006),⁵⁸ but I see a major problem with defining attention as mere accessibility. It would mean that every dormant, potentially accessible, yet actually unaccessed trace in long-term memory is "attended" merely because it is potentially accessible. That seems to make attention a trivial and useless concept, and bears little similarity to how attention is used in the literature.

⁵⁷ See also attention-as-selection-for-action below (3.3.4.2.4).

⁵⁸ Chalmers considers availability as defining a "modified notion" of *access consciousness*, which I argue is roughly equivalent to Executive Attention (3.4.4 and 7.2.2) —the kind of *attention* that is of interest for Q. Mogensen and Overgaard (2018) repeatedly assume that *phenomenal consciousness* is "availability." Seager (2016, pp. 203–204) critiques Prinz's assertion that consciousness is accessibility to working memory using the "meddler device" objection.

Neither does Block's (2007) Aristotelian distinction between "actually accessible" and "potentially accessible" solve this problem. His distinction points out that while a vast amount of content is always *potentially* accessible by attention, the capacity limitations of attention (8.3) severely limit the amount of content that can be *actually* accessed at any given moment. But notice here that the concept doing all the work of defining content as attended—the substantial defining feature—is still the actual access, not the accessibility. Let us say there are a thousand potentially accessible traces in long-term memory, but only any five of them are actually accessible at any given moment. If we define "attended" as "those traces that are actually accessible," how do we determine which traces are attended? By observing which traces have actually been accessed. For the purposes of defining (or at least identifying) attention, "actually accessible" plays *exactly the same role* as "actually accessed," even if it has other useful roles to play in other discourses. Thus, potential accessibility fails to define attention, and actual accessibility just collapses into actual access. Accessibility does not define (or identify) attention.

3.3.4.2.2 *Detection and Recognition*

Another approach focuses on the idea of *detection* or *recognition*. Attention is the process of "detecting signals⁵⁹ for (conscious) processing" (Posner & Petersen, 1990, p. 26). This process of detection can itself be a complex one, and we can identify a number of possible stages within it in which attention may be invoked. These include: *orienting*; *vigilance* or *monitoring*; *scanning* or *searching*; *expectation*; *alerting*; and *tracking*.

Orienting is the act of rearranging one's physical deportment in space in order to better receive information from the environment. Classical examples of this are turning one's head to see better, and rapid saccades of the eyes. Petersen and Posner (2012) consider orienting to be one of three defining characteristics of attention,

⁵⁹ I imagine the signals may be sensory inputs or abstract thoughts, emotions, beliefs, etc.

together with alerting and executive control. But Prinz considers orientation and attention to play two distinct roles: “Informally, orienting alters what information gets in, and attention alters where it flows” (Prinz, 2011, pp. 193–194). Nonetheless, the act of orienting itself—whether physically moving body parts or even the shifting of the focus of spatial attention while remaining physically still—does seem to constitute an integral part of some kinds of attention. It is harder for the gorilla to remain unattended if one’s eyes orient and fixate upon it, and the lack of orientation among the subjects is a reliable indicator of the absence of attention to the gorilla.

Vigilance or *monitoring* is non-specific receptiveness to *any* new incoming content (contrasted with *searching* for a specific Target, below). The gorilla would have been attended by subjects had they been asked to simply report what they saw, rather than count the number of passes. This state of expectancy has been called a kind of “preparatory attention” (Zeman, 2001, p. 1274).

The pass counting task itself was an example of the more narrowly specified *scanning* or *searching*—the goal-directed search for something in particular. This is James’ (1890, p. 434) idea that one of the functions of attention is to hold in the mind an image of the thing one is searching for, so that one can identify that thing by comparison to the image.

Closely related to monitoring and scanning is Clark’s (2016, p. 62) concept of *expectation*, a particular mode of searching where the subject is cued to expect a stimulus at a certain spatial location, and for which he presents a predictive processing account. The gorilla is not attended because it is not expected.

The act of finding the object of the search is what Petersen and Posner (2012) call *alerting*—the signal that something has actually been detected or recognised. This alerting is absent in the case of subjects who fail to see the gorilla, and reliably indicates the absence of attention to the gorilla.

Tracking is what happens when one has found a Target, and then maintains focus on that Target over a period of time⁶⁰. The subjects identified the basketball, attended to it due to a top-down imperative, and maintained that attention as the ball was passed from player to player. O'Regan and Noë (2001, p. 944) suggest that tracking a target while otherwise occupied is the signature of moving from inattention to that target, to attention. In 8.3.2 I consider divided attention in tracking multiple moving objects.

3.3.4.2.3 *Selection*

Our cognitive economy is engulfed constantly in a raging flood of information, whether from outside, via perception, or from within. How can it deal efficiently with all this information and tame it into James' efficient stream of consciousness? How can it sift the relevant from the irrelevant, and choose to act accordingly? One of the most prominent definitional approaches to attention views it as being primarily *selection*,⁶¹ a selection that allows the mind to efficiently manage this constant flood of information. In the mid twentieth century, the *bottleneck* model of attention enjoyed a vogue, where attention was seen as a unitary mechanism of constriction of processing capacity at some stage in the chain of cognitive processing (Broadbent, 1958, 1971; Mole, 2013, sec. 1.6; Smithies, 2011, pp. 250–251). This led to the question of whether such selection occurs early or late in the processing chain (Pashler, 1998, pp. 13–19), which in turn led Allport to challenge the unitary nature of consciousness. More recently, Baars (1997b, p. 368) characterises attention as *selection* in contrast to consciousness as *experience*, while Campbell (2011) contrasts *selection* with *access*, and argues that a constitutive relationship between attention and consciousness may be found in selection, but not in access. I describe a variety of ways in which the general idea of selection has been applied to the

⁶⁰ For an evolutionary account of the tools humans use to make their tracking more effective, see Menary (2018).

⁶¹ "The selective nature of attention has perhaps been the most widely studied area" (Rees, 1999).

definition of attention: *filtering*; metaphors of the *gateway*, the *spotlight*, and *zooming in/out*; and the ideas of *exclusion*, *competition*, and *selection for action*.

At the heart of the bottleneck model is the idea of *filtering* (Broadbent, 1958)—reducing a large and unruly amount of information to a smaller, more manageable amount through a process that selects the information that is more likely to be relevant and useful to the subject’s current goals and needs (the gorilla, being irrelevant to the present task of counting passes, was filtered out of the subject’s focus). The *frame problem* in artificial intelligence⁶² can be seen as being solved in human cognition through attention-as-filtering. Using context from past experience and current contextual hints from present perception, attention filters out the relevant and useful content for the further processing that will produce results that achieve the subject’s goals. However, the nature of this filtering is complex, as is its relationship to consciousness (Baars, 1988, pp. 34–36).

Attention has been defined by what it selects *for*—the metaphor of a *gateway* leading to something else (further down the processing chain). Examples of attention-as-gateway include the gateway to: visual processing (Desimone & Moran, 1985); consciousness (Baars, 1988, p. 369; Crick & Koch, 1990, p. 269; Mack & Rock, 1998, p. 25); and working memory (Awh, Vogel, & Oh, 2006, p. 202; De Brigard & Prinz, 2010, p. 52). Of course, these gateways are not mutually exclusive—as an umbrella term, attention may be at once selection for visual processing, consciousness, working memory, and any number of other processes.

Spotlight theories of attention (C. W. Eriksen & Hoffman, 1972) emphasise the idea that what determines whether or not attention is being paid to a stimulus is its location (Mole, 2013, sec. 2.7). While spatial location does seem often to be a major

⁶² The frame problem first arose as a technical problem in artificial intelligence (McCarthy & Hayes, 1969) where it has largely been resolved today. Roughly, it is the problem about how to know what to sift out from the flood of data available in a real-world environment so that only what is relevant to the task at hand is focused on. It continues to be debated and explored in cognitive science (Dennett, 2006; Shanahan, 2016).

factor in how attention is directed, there are good reasons for holding that it is certainly not the *only* factor. For example, one might argue that the gorilla undoubtedly entered the spotlight of the subjects' attention, as when the ball passed directly in front or behind it, yet it still failed to be attended. Pure spotlight theories cannot account for pop-out in visual search where what brings the stimulus to attention is its contrast against a background, not its spatial location. Neither can spatial attention explain olfactory attention, or much of auditory attention, which are not spatial by nature.⁶³

The metaphor of a spotlight naturally raises the question of the dimensions of that spotlight. How narrow or wide is its beam? And is it possible to vary its dimensions, to “*zoom in/out*” as it were (C. W. Eriksen & St James, 1986)? The gorilla is not attended because the spotlight has been zoomed in very tightly upon the ball, so that in fact, the gorilla never is in the spotlight, or at least, never is the whole gorilla in the spotlight all at once, and perhaps never is it in the spotlight long enough, to create a powerful enough bottom-up signal to draw the subjects' attention to it.

Selection involves the inclusion of some content, which therefore means that other content is *excluded*. When one selects the basketball, one is automatically thereby not selecting (excluding) the items left over, including the gorilla. This approach to defining attention may be traced back at least to JS Mill: “The expression [attention] means that a sensation tends more or less strongly to exclude from consciousness all other sensations” (cited in De Brigard & Prinz, 2010, p. 52) and Treisman (2003, p. 102) contrasts the metaphor of an actively excluding attentional *window* with that of an inclusive attentional spotlight. Thus, exclusion may occur via *suppression*. van Swinderen (2005, pp. 327–328) suggests that in both sleep and alert attention (as opposed to being in an alert yet inattentive state), we raise the threshold required by peripheral inputs to enter into consciousness.

⁶³ Keller (2011) argues that the current focus on visual attention at the expense of other modalities risks obscuring important aspects of attention.

The intense concentration on the attended target to the exclusion of all other signals may be described as *absorption* (Tellegen & Waller, 2008). An extreme case may be savants like Kim Peek, where a suppression of distracting stimuli seems responsible for their superhuman powers (Treffert & Christensen, 2005). Related to absorption are the total *dedication* of available resources to a task at the expense of other tasks (Mole, 2011b, p. 67) and the idea of *absence of distraction* (e.g., Cowan & Morey, 2006; Engle, Kane, & Tuholski, 1999, p. 104).

One way that selection of cognitive content can occur is via a process of cognitive *competition*. Thus, in the competition for cognitive resources, the basketball wins hands down over the gorilla, due to the top-down imperative to focus on counting passes. Mole (2013, sec. 2.6) considers competition models to be the “clearest non-bottleneck mechanisms for achieving selectivity,” and discusses the sub-varieties and nuances of competition models. One of the most influential of these is *biased competition* (Desimone & Duncan, 1995; Ruff, 2011), which introduces the idea that other processes in the brain, whether top-down, bottom-up, or lateral, interact with the process of attention to bias its selection “choices.” Some however, have sought to distinguish competition from attention, arguing that “attention doesn’t refer to competition, as such, but, rather to a process that occurs when a competition is won” (Prinz, 2011, p. 183).

Finally, recognising that attention is not restricted to perception is the idea of attention as *selection for action* (Allport, 1987; W. Wu, 2011, 2016).⁶⁴ “Action” is used here in the more general sense that encompasses purely cognitive activity (thinking, remembering, calculating) as well as motor activity, whether physically performed or merely mentally rehearsed or intended, so this definition is closely related to attention-as-influence (3.3.4.2.1). Subjects selectively attended only to content relevant to the action of counting passes. Others have also seen an intimate connection, perhaps even a unity of mechanism, between attention, working memory and motor action (Postle, 2006; Theeuwes, Olivers, & Chizk, 2005).

⁶⁴ This was one of Treisman’s attentional strategies (3.2.1).

3.3.4.2.4 *Control*

Human goals tend to be complex and extended over time. Coordinating the cognitive processes required to sustain the pursuit of such goals⁶⁵ is therefore a complex task. When the subjects set out to count the number of passes, one might argue that they could not achieve this goal without *attention-as-control*. Parasuraman (1998, pp. 7–8) points out that such mechanisms of attentional control do indeed feature heavily in executive and planning components (D. Norman & Shallice, 1986) of popular models of working memory (Chapter 5). Attention can therefore be thought of as the mechanism that directs the cognitive traffic in a cognitive economy in order to ensure the smooth and most effective and efficient functioning of that economy in its constant striving to fulfil those complex goals.

3.3.4.2.5 *Coherence*

If attention serves to make content more manageable, one of the ways it might achieve this is by making it *coherent*. On these accounts, attention may be defined by its functional role in separating out and coordinating a limited, relevant, and coherent body of information as the foundation for further processing in things like reasoning, agency, and action. The sensory data from the basketball passes is processed to form a coherent story about the movement of the ball from player to player in the accepted context of a basketball game (of sorts). The trespassing gorilla is incoherent—difficult to reconcile with this context, and so is ignored, i.e., unattended. Attention-as-coherence features in the *Feature Integration Theory* (Treisman, 2003; Treisman & Gelade, 1980), where attention helps bind disparate features into coherent perceptual wholes. In predictive coding models, attention is the process of optimising precision of signals, a kind of coherence (Friston, 2009;

⁶⁵ Compare Ganeri's (2017, p. 225) concept of *intending*, which is more the Phenomenal striving towards a goal.

Hohwy, 2012). Mole's (2011a, 2011b) *cognitive unison view* also plausibly falls under this subcategory. In a similar vein, Wyble (2015) creatively describes visual attention as being not so much a filter or a newspaper editor, picking which stories to publish, as like a movie editor, arranging scenes into a coherent story.

3.3.4.3 *Organisation*

The final category of Bechtel's framework focuses not on which parts of the brain are involved in attention, nor on what those parts do, but on how the parts are organised in relation to each other. In relation to attention, this *organisation* can be examined on many levels: the relationships of individual neurons to each other; of networks of neurones to each other, of regions to each other, or of functionally specific connectomes to each other. Alternatively, we could look at organisation at Marr's (1982) three levels of computational theory, representation and algorithm, and implementation. Organisation is not merely spatial but also temporal. Of course, these are all not mutually exclusive accounts, but complement each other. A thorough analysis of attentional organisation is far beyond the scope of this thesis and will not play a major role in identifying the presence or absence of attention,⁶⁶ since it is as difficult to identify in living subjects as parts are (3.3.4.1), so I will content myself with some brief remarks relevant to my arguments.

In what sense might organisation define attention? Bechtel's triad of parts, operations, and organisation are meant to be intimately connected. The operations that define attention (3.3.4.2) will not only require the right parts in the cognitive system, but also that those parts be organised in certain ways that allow for those operations. In other words, there may be certain patterns of organisation that are infallibly correlated with attention and might therefore serve to identify its presence or absence.

⁶⁶ It bears only indirectly on the discussion of neural signatures and no-report paradigms (7.4.3).

There is some hope that we might one day be able to identify the organisation of the various subtypes of attention, and how they relate to each other, whether neuronally or functionally. For example, Rosenberg et al., (2017, p. 299, Box 1), adopting an individual differences approach, found that some of the functions subsumed under the umbrella term, “attention” vary among individuals together, whereas others vary independently of each other. Thus, the ability to sustain attention over long periods of time does not necessarily covary with the ability to multitask. So also, spatial orienting, attentional capture, and inhibition of return seem to vary independently of other sub-functions. On the other hand, functions like search, tracking, and visual short-term memory all seem to vary together across individuals and over time, suggesting they may depend on some kind of common attention factor, at least in part. These patterns of co-variance might indicate underlying organisational patterns.

It will come as no surprise that the organisation characteristic of attention is complex. Attentional Operations require content to relate to other content, and processes to influence other processes. We should expect them, therefore, to be characterised by organisational patterns such as *cyclic pathways* and *feedback loops* (Bechtel, 2008, p. 17). What is more, parts and operations involved in attention will be intimately entwined—richly *contextually interconnected*—with other cognitive systems, contentfully, functionally, neuronally, and in terms of organisation. This has led some to identify organisational neural structures that allow for recurrent neural firing patterns with attention (Ruff, 2011, pp. 8–9), although others have identified them with consciousness (Lamme, 2010).

This kind of interconnection will be significant when we come, for example, to explore attention in the visual system in Chapter 7 or tease apart iconic memory from working memory in Chapter 8. However, while highly interconnected organisational patterns may characterise attention, they are also nearly ubiquitous throughout all cognitive processes, so they will hardly do to identify attention or the lack of it. Whether there are organisational patterns that are unique and specific to attention remains to be seen.

In summary, then, Behaviourist and Phenomenal definitions both beg the question of Q, albeit in different ways. Among Mechanistic definitions, defining attention by virtue of its Parts or Organisation is fraught with difficulties. However, functional or Operations definitions of attention are much better developed and much more useful. It is therefore the Operations definitions of attention upon which I chiefly rely in the arguments that follow. In this I will not be alone, as it has been my observation that most discussions of attention and consciousness assume—implicitly or explicitly—one or more Operations definitions of attention.

3.3.5 Is It Really Attention *per se*?

It could be argued that some of the definitions in my taxonomy may be *precursors* to attention (e.g., monitoring, vigilance, salience) or *consequences* of attention (influence, alerting, control, coherence), but are not rightly classified as attention *per se*. There are two responses to this objection.

First, my purpose here has not been ontological, but practical. The purpose of my taxonomy is to provide a firm basis for identifying the presence or absence of attention in cognition. Whether the Operations definitions are constitutive or attention or merely infallible markers of attention is immaterial to this purpose. Of course, the infallibility is crucial, but I believe there is a way to deal with it (3.4.4).

Second, I believe that there is a unifying principle that plausibly captures and unifies the Operations definitions of attention in my taxonomy essence. This principle warrants the inclusion of all the definitions above in the taxonomy, but also provides a practical and effective basis for discussions of attention like Q. I develop this argument in the next section.

3.4 My Approach to Defining Attention

The purpose of this conceptual analysis of attention is to develop an approach to defining attention that is optimal for addressing Q. One way of doing that might be to look at all the different varieties of Operations above and ask what it is that ties them together. Why have they fallen under the “umbrella” of attention? What qualifies or disqualifies something to be called attention? In this section, I propose that attention is a suite of strategies for structuring cognition, and suggest the essence of attention that ties all the Operations definitions together is the broad idea of selection for further processing. I then provide some examples of other cognitive strategies that do not fall under this umbrella. I further suggest principles for maximising the utility of my approach, and to that end, draw a final distinction between Liberal and Executive Attention.

3.4.1 The Essence of Attention

What kind of thing is attention? It has been suggested that attention is a system (Posner, 1994, p. 7399); a process (Kentridge & Heywood, 2001); and a variable (Baars, 1997b) or an adverb (Mole, 2011a) of cognitive processes. But the Operations definitions above are best captured by another possibility—attention is the *structuring* of cognition (Watzl, 2010, 2011a). The Operations enumerated above are, I submit, best thought of as *strategies*, indeed, a *suite of strategies for structuring cognition*. A “strategy” account avoids the problems inherent in reductionist accounts of attention (Watzl, 2011b, pp. 846–8). It is thus intimately related to

systems and processes that implement those strategies, which may therefore be said adverbially to be attentional just when they do.⁶⁷ The essence of each of the Operations definitions is metaphysically best described as a multiply realisable strategy (Bickle, 2016)—even security cameras suitably equipped with servo motors can monitor, orient, and track. And the apparent lack of an “attention centre” in the brain suggests that attentional strategies are implemented throughout cognition (Watzl, 2011b, p. 847).

What is it that ties the Operational definitions of attention above together, and distinguishes them from other cognitive strategies? What captures the essence of each definition? A very good candidate is the concept of *selection for further processing*.⁶⁸ The “further processing” can be anything: storage; motor action; binding; reporting; etc. This makes perfect sense of the *Selection* definitions, but the other Operations definitions capture nuances of the selection for further processing that occurs in human cognition. Thus, without *Access* to content, it cannot be further processed, and it is just that content that is accessed that is further processed. *Detecting* and *Recognising* content is sometimes the first step in that selection for further processing, without which it could not proceed. *Control* is necessary to direct the selection and further processing to the desired goal. And *Coherence* ensure efficient structuring for further processing that fulfils goals. Operational attentional strategies are thus intimately interconnected in a constant dance with each other, and with perception, motivation, motor command, and consciousness (Merker, 2007). The working definition of attention that I use in this treatise is thus: *a suite of strategies for structuring cognition for further processing*.

⁶⁷ This is why I rejected “accessibility” as an attentional definition above (3.3.4.2.1). We can apply the adverb “attended” iff one of the attention strategies is actually being implemented, not if it is only potentially implementable.

⁶⁸ Compare Jennings’ (2012, p. 536) Phenomenal definition of attention as “the act of mental selection.” Her emphasis is on the subjective phenomenal *act* as opposed to an objective *event* of attention, both of which I eschew in favour of an objective *strategy*. Fazekas and Nanay (2018) propose an alternative unifying principle to selection—amplification—which I subsume under my *Access* category, and which does not capture the many other Operations.

3.4.2 Non-Attentional Strategies

This understanding of attention will only be useful though, if attention is not just a term for any kind of cognitive structuring strategy at all. There are other strategies for structuring cognition that are not strategies for *selection for further processing*, but are either pre-conditions for attentional structuring or constitute the further processing itself. For example, *storage* is conceptually dissectible from attention. The strategy of encoding content in a retrievable way requires selection of that content *for* encoding, but encoding is not itself selection. The essence of storage is that idea of temporal endurance, not selection. In a similar way, the essence of *binding* is not selection, but combining content in principled ways.⁶⁹ Attention-as-coherence is a strategy implemented to help achieve that binding in useful ways, but is conceptually dissectible from the combining itself. Similar dissections can be made for other strategies such as calculation, retrieval, comparison, and many more.

3.4.3 Using “Attention”

On this account, attention is present when and where *any* one or more of the Operational attentional strategies is being implemented, and it is absent when *none* those strategies is being implemented. The various distinctions above (3.2.3) mean there is ample room to move, however. For example, there may be *spatial* attention-as-selection in the absence of *object* attention-as-selection. And if attention is selection for further processing, the reference point of “further” needs to be identified before we can determine whether attention is present or not. For

⁶⁹ *Pace* Triesman and Gelade (1980).

example, when I attend to barely legible script on a piece of paper, there is a particular implementation of attention that transforms my perception of that script from illegible to legible. That is not to say attention was not implemented in my perception of the script when it was still illegible, but the two attentions have different reference points beyond which the “further” of “selection for further processing” applies.

I have taken a descriptive approach to attention, but there remains a proscriptive question: how should we use the term “attention?” My approach is in no way meant to be proscriptive except in the broadest terms. It is perfectly valid to use just one or a few Operational definitions, for a particular narrow purpose, so long as one is aware that that is what one is doing. A study on attention-as-tracking need not bother with attention-as-competition, although in most implementations of attention, we do indeed find many definitions co-occurring and overlapping. Human beings tracking moving objects implement not only tracking, but orientation, selection-for-action, suppression of other stimuli, and so on. However, when it comes to broader questions like Q, the whole suite of strategies is in play, particularly when we want to identify a complete absence of attention.

3.4.4 Liberal and Executive Attention

The virtual ubiquity of attentional strategies in cognition (Serences & Kastner, 2014) threatens to render the concept impotent when it comes to Q. If content must be attentionally selected in order to be processed, and we become conscious only of content that has in some way been processed, then surely there can be no consciousness without attention. Another distinction saves us from this dilemma. Let us call any implementation of attentional strategies whatsoever—at any level of cognitive processing—a liberal definition of attention, or *Liberal Attention* for short. And let us distinguish this from attentional strategies implemented specifically by the cognitive executive (I flesh this idea out in Chapter 5 when I discuss working

memory)—*Executive Attention*.⁷⁰ Now it is clear that not all attention is Executive Attention. Attention in the brainstem systems that maintain the body's homeostasis (blood pressure, body temperature, etc.) is certainly not Executive Attention. But the attention implemented to visually track basketball passes is Executive Attention. Note that the defining feature of Executive Attention is that it is implemented by the executive. As such, it can cut across all those distinctions drawn in 3.2.3.⁷¹

This distinction provides a principled way to pose Q in a substantive and interesting manner: *is Executive Attention both necessary and sufficient for phenomenal consciousness?* I am aware that this is indeed what most authors on the question have in mind, but occasionally there can be some confusion between Liberal and Executive Attention that can muddy the waters, so it is worth making this distinction clear. What is more, I observed above (3.3.5) that all that is needed for the purposes of Q is for the Operational definitions to be infallible indicators of attention. So long as we can be sure that an attentional strategy is being implemented by the cognitive executive, that infallibly constitutes Executive Attention, as I have defined it. And so long as we can be sure that the executive is not implementing any of the attentional strategies with regard to the Target in question, that infallibly constitutes an absence of Executive Attention. Executive Attention—properly defined—is thus a potent tool for exploring Q.

⁷⁰ Note that this is not the Personal/Subpersonal attention distinction I discussed above. That distinction was about reducing Personal level attention to Subpersonal. Here, Executive Attention is implemented by the coordinating system of cognition, which may receive content produced by other specialised systems implementing Liberal attention, but is distinct from them. My concept of Executive Attention is also broader than that Kane and Engle's (2002, p. 638) concept of "executive attention," which captures only some of the strategies in my Taxonomy.

⁷¹ My approach also meets JH Taylor's (2015a, pp. 40–41) two criteria for defining attention. It is *extensionally adequate* in that it allows for varying degrees of "leniency" in defining attention, depending on the needs of the task at hand. It is *non-circular* in that it is based on a descriptive approach rather than a stipulative approach.

3.5 Chapter Summary

After briefly summarising some pertinent distinctions and difficulties in defining what attention is, I adopted a descriptive approach and surveyed the literature to develop a taxonomy of ways attention has been defined that will inform the discussions to follow. From the Operations class of this taxonomy I distilled a working definition of attention as a *suite of strategies for structuring cognition* that have in common the concept of *selection for further processing*. The presence of any one or more of these strategies is the criterion by which we may determine that attention is present, and the absence of them all identifies the absence of attention. Finally, I argued that in addressing Q, Liberal Attention is uninteresting, since it is ubiquitous in cognition. The really interesting question is the relationship between Executive Attention and phenomenal consciousness. I turn to the nature of that relationship in the next chapter.

4 Relationships

4.1 Introduction

There are two big issues that have muddled the waters somewhat in the discourse over Q. Definitional clarity of the terms involved is one, which I addressed in Chapters 2 and 3. The other is incomplete, imprecise, or suboptimal frameworks for analysing the possible relationships between attention and consciousness. It is often difficult to ascertain what exactly an author's position is on Q, or to relate that position to those of others. In this chapter, I attempt to make some progress in bringing clarity to this area. I develop a framework for addressing Q that can not only relate the many different approaches taken by diverse authors to each other, but perhaps even be applied to questions other than Q, and even beyond cognitive science. While this chapter represents part of my literature review, it also aims to make an original contribution to study of relationships between any two concepts in any field.

I preface this chapter with a brief look at the history of thinking about the relationship between attention and consciousness, as an introduction to an overview of the discourse on this question over the past thirty years. I then introduce and develop a conceptual framework that goes some way to alleviating the problems inherent in the discourse to date. I compare it to some of the existing (mostly implicit) frameworks in the literature, before employing it to bring some order to the various views of the chief authors involved in this discourse. Patterns of relationship do not necessarily disclose the reasons for their obtaining, so finally, I describe four possible underlying reasons for the relationship between attention and consciousness.

4.1.1 History

The relationship between what we today call attention and consciousness is not just a modern issue. For example, the Stoic Chrysippus observed that both humans and animals perceive far more than they attend to (A. Long, 1986, pp. 172–173). He developed a theory of how both choose what to attend to and what to ignore, based on their natural disposition and their relationship to their environment. This process, governed by “impulse” in animals and infants, comes under the command of reason in mature humans: “reason supervenes as the craftsman of impulse.”⁷² “Reason” here may plausibly be taken to denote what we today would call “conscious volition” or “self-regulation,” as Petersen and Posner (2012, pp. 82–84) put it in their eerily similar (though thoroughly modern and scientific) treatment of the same topic.⁷³ Of course, Chrysippus’ main interest was in determining how reason uses attention to live a virtuous life, but it is evident that it is at least part of

⁷² Diogenes Laertius, VII, 86.

⁷³ Compare Chrysippus’ maxim above to Petersen and Posner’s “Both behavioral and resting state functional data suggest substantial development of the executive attention network between infancy and childhood” (page 84).

this project to attempt to characterise the relationship between attention and conscious cognition.

In ancient *Abhidharma* Buddhist philosophy, the concept of *manasikāra* or attention, is considered to be one of the “five omnipresent factors” (Dreyfus & Thompson, 2007, pp. 98–100). This would seem to prefigure those in the modern debate who consider attention to be “omnipresent” or coextensive with consciousness. However, the tradition distinguishes *manasikāra* from two other factors: *samādhi* or concentration on the object of thought; and *smṛti* or mindfulness, the ability to remain focused on the object of thought without wandering. Unlike *manasikāra* neither of these factors are considered to be “omnipresent.” These distinctions are antecedents of some of the different definitions of attention discussed in Chapter 3, but again, we see the intuitive appeal of relating forms of attention to conscious experience.

Similar discussions may be found in many other philosophical traditions, such as (but certainly not limited to): the concept of contemplative (*θεωρία*) prayer in the Christian tradition and its roots in Platonic mysticism (Louth, 2007); Augustine’s view that attention plays a role in generating experience (Normore, 2016); Avicenna’s discussion of self-awareness in the Islamic tradition (Kaukua, 2015); the Scholastic discourse on the requirements for self-knowledge (Perler, 2017); and Leibniz’s view that rationality requires reflection—turning one’s attention to one’s own conscious thoughts (Bender, 2016).

4.1.2 Modern Discourse

Modern discourse on the relationship between attention and consciousness is often considered to really take off with William James, who considered attention to be a particular aspect of conscious experience (James, 1890, pp. 403–405). Wilhelm Wundt saw attention as the choosing from among that which was already conscious

(Wundt, 1912, p. 16). But Phillips (2011a, p. 221, note 4) observes that “discussion of attention and consciousness in this period is unsatisfactory owing to deep disagreement over the nature of attention.”⁷⁴

“When the disciplines of philosophy and psychology split from one another in the first half of the twentieth century, it was psychology that got custody of attention while philosophy was given responsibility for consciousness ... When the discipline of cognitive psychology was getting under way at the end of the 1950s, nobody felt much need for a definition of *attention*. Research into the basis of attention could proceed without such a definition because there was almost universal agreement about which experimental tasks involve attention” (Mole, Smithies, & Wu, 2011b, p. xi).

The lack of interest in consciousness that prevailed until the second half of the twentieth century meant, of course, that there was a corresponding lack of interest in the relationship between attention and consciousness. It is only toward the end of the twentieth century that we begin to see seminal papers stirring debate and discussions. Block’s (1995) much discussed distinction between access consciousness and phenomenal consciousness provided the vocabulary and inspiration for a rich discourse. Hardcastle (1997) challenges authors from the previous ten years who assert that attention and consciousness amount to the same thing, and that William James took them to be so, arguing instead that the two are distinct phenomena. Hardcastle’s view exemplifies the empirical flavour of recent attempts to address Q, as does other influential work from the 1990s (e.g., Kentridge, Heywood, & Weiskrantz, 1999). Around the turn of the millennium, discussion of this relationship begins to really take off, with an increasing

⁷⁴ See Phillips’ pages 204–205 for a succinct account of the discourse around the turn of the twentieth century and Combs et al., (2010) for a more detailed account.

convergence between the disciplines of experimental psychology, philosophy, and neurobiology. This recent discourse is surveyed in 4.3 below.

4.2 The Set Theoretical Framework (STF)

“inquiry proceeds conceptually and empirically in tandem to uncover the most fruitful ways of delineating the subject matter”
(Gross, 2018, p. 2).

Authors on Q present a bewildering variety of approaches, definitions, and concepts, which has led to “the baffling emergence of parallel literatures, akin to divided universes that reflect one another, but scarcely interact” (Tong, 2013, p. 489). It would be very helpful to have a framework against which all these divided universes could indeed interact, but no such framework exists for Q. In this section, I invent one.

In brief: I develop a descriptive framework that exhaustively defines the conceptual space in which this discourse is conducted. I take a subject’s cognitive economy—the sum of all their mental life—to be comprised of particular cognitions. Each of these cognitions may or may not be attended and/or conscious. This gives us four possible *Combinations* that describe the patterns of absence or presence of attention and consciousness in any given *particular* cognition. When we come to describe the cognitive economy *as a whole*, we must take into account which of these Combinations are instantiated in the cognitions that comprise that economy, and which are not. I call a pattern of Combinations in a cognitive economy a *Scenario*. Thus, for example, there is a Scenario where some of a subject’s cognitions are

neither attended nor conscious, some are both attended and conscious, some are attended but not consciousness, but none are ever conscious but not attended. Since particular cognitions are considered as members of sets (i.e., Combinations) and since I use the language of set theory, I have called this system the *Set Theoretic Framework* (STF).

After deriving and explaining my framework of Combinations and Scenarios below, I highlight the advantages of STF over other comparable frameworks or taxonomies, and then use it to frame a literature review of recent work on Q in 4.3, which in turn provides the springboard to my own arguments in Part II. STF promises to provide a common interdisciplinary framework for discourse about Q.

4.2.1 Combinations

A particular cognition—e.g. seeing a red apple—may be described with respect to attention and consciousness by one of four logically possible *Combinations*,⁷⁵ described just below. A whole cognitive economy, comprised of many particular cognitive processes or cognitions, each with its own Combination, may thus be described by one of sixteen logically possible *Scenarios*, described in the following section.

Before I begin to develop the concept of Combinations, three preliminary remarks are necessary to bracket certain rabbit-holes I wish to avoid falling into. First, the Combinations described below describe only *patterns* of correlation. They are metaphysically neutral as to what attention or consciousness are, and as to the *reasons* for the relationship between them (which I address in 4.4). STF makes no *a*

⁷⁵ I use “Combinations” rather than “Permutations” as the order of items matters in the latter, but not in the former.

priori assumptions that attention and consciousness are identical or perfectly co-occurrent. For example, while Phenomenal definitions of attention (3.3.3) can be expressed in the language of STF, they are just one of the many possibilities captured by STF. The exhaustive nature of STF allows us to address Q without begging the question of their relationship (3.2.4.2).

Second, for the purposes of this discussion, I have preferred the terms “cognitive processes” or “cognitions” over “cognitive states” as the latter has connotations of a static synchronic snapshot, whereas the others imply more of a diachronic event, which is an important aspect of the nature of the cognitions we are considering.⁷⁶ “Processes” or “cognitions”⁷⁷ thus also seem less confusing than terms like “brain activity states” (Dehaene et al., 2006, p. 204) that combine the static (“state”) with the dynamic (“activity”).

Third, while the Combinations below represent the relationships between just two particular aspects of cognitions (i.e., attention and consciousness), it is important to always bear in mind that any cognition—and indeed, the overall cognitive economy—is immensely complex, and profoundly interconnected. Cognitions almost never occur in a vacuum.

If we assume, then, that attention and consciousness are possible descriptions of cognitive processes or cognitions, there are exactly four logically possible ways to characterise any given cognition with respect to attention and consciousness—four *Combinations*.⁷⁸

⁷⁶ For an in-depth discussion of these kinds of issues, see Soteriou (2013) and Mole (2016).

⁷⁷ In general, I use “cognitive process” where the physical brain process is more pertinent, and “cognition” where the mental perception or thought is more pertinent.

⁷⁸ The examples given for each Combination are illustrative only and open to different interpretations, so they are not to be taken here as arguments that the Combination is instantiated. Those arguments are in Part II.

Combination 1. ~A~C: neither attended (A) or conscious (C)

- A particular cognition may lack both attention and consciousness.
The subject may be attending to something else, or not attending at all, and she may be conscious of something else, or not conscious at all.
- This Combination is relatively uncontroversial. That is to say, most people would accept that it occurs, at least sometimes.
- An example of this Combination might include the “housekeeping” cognitive processes going on in the brain while one is in a deep, dreamless sleep (Lewis, 2013).

Combination 2. A~C: attended, but not conscious.

- A particular cognition may involve attention, but not consciousness.
- This Combination is more controversial, with some people arguing that it never occurs.
- An example of this Combination might be a cue in the blind region of a blindsight patient’s visual field that, despite being phenomenally unconscious to the subject, nonetheless directs her attention to a second region (Kentridge et al., 1999).

Combination 3. C~A: Conscious, but not attended.

- A particular cognition may involve consciousness, but not attention.
- This Combination too is controversial.
- An example of this Combination might be the famous Sperling (1960) experiment in which all the letters in a grid seem to be consciously experienced, but only a small subset of them can be attended and subsequently reported.

Combination 4. A&C: Both attended and conscious.

- A particular cognition may involve both attention and consciousness.
You consciously experience the object of attention and attend to the object of conscious experience.
- This Combination is almost unanimously uncontroversial.
- It comprises a great deal of our common daily experience, such as when you pleasurably attend to your phenomenal experience of the colour of a rose and the particular shade of its redness.

Much of the contemporary discourse about Q consists of making a case for the instantiation (or lack thereof) of one or more Combinations in human cognitive economies. In Part II, I take Combinations 1 and 4 to be uncontroversial, briefly catalogue the powerful evidence for Combination 2, and spend the balance of this treatise carefully exploring the evidence for Combination 3: *C~A*, for which I argue there is a very good case to be made. While *deriving* an answer to Q requires exploring each Combination, *expressing* an answer to Q requires us to bring together the answers about whether or not all four Combinations are instantiated or not. To do this, we need to describe *Scenarios*⁷⁹—patterns of occurrence of Combinations in a whole cognitive economy (which is what Q is really about).

⁷⁹ In 6.2 I discuss the possibility of *global* Combinations and distinguish them from Scenarios.

4.2.2 Introducing Set Theory

Scenarios are derived by applying standard set theory to the Combinations of attention and consciousness in a cognitive economy. We can think of attention, as being a set, A , the members of which are all particular instantiations of human cognitions in the cognitive economy of an individual subject, that are characterised by involving attention. We can think of consciousness, C , in the same way.⁸⁰ We can thus represent the instantiations of attended and conscious human cognitions using the Venn diagrams of standard set theory⁸¹ as shown in Figure 1.

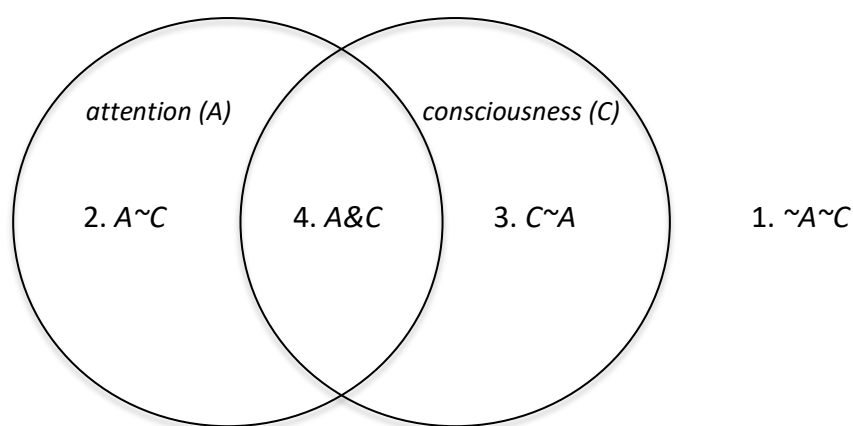


Figure 1 The four possible *Combinations* of attention (A) and consciousness (C).

It needs only one irrefutable empirical instance of a Combination to confirm that that Combination is instantiated in subjects. On the other hand, the absence of even one irrefutable instance of a Combination, while highly suggestive, is not proof positive that it can *never* be instantiated—it is impossible to prove a negative

⁸⁰ More formally, the extension of the property of involving attention is the set A , and the extension of the property of involving consciousness is the set C .

⁸¹ For set theory, see Smith (2012, Chapter 16).

(absence of evidence is not evidence of absence). For example: we may simply not yet have stumbled across that Combination empirically; or we may not yet have recognised it as such; or it may not be instantiated in normal human cognition, but it may still be instantiated in abnormal human cognition; or in actual non-human cognition (e.g. in dogs); or in hypothetical or fictional non-human cognition (e.g. Martians or more evolved human beings of the future). I address this question in 4.2.3.2.

4.2.3 Scenarios

While only one Combination can describe any particular *cognition*, any, all, or none of the Combinations may be found together within a whole *cognitive economy*, which after all, is composed of numerous individual cognitions. A *Scenario* describes the pattern of Combinations that are instantiated for any given subject (whether synchronically or diachronically). It turns out that there are sixteen logically possible Scenarios that I describe below. These sixteen Scenarios exhaust the possible positions on Q.

For a particular cognition, Combinations are mutually exclusive (only one Combination can describe a particular cognition), but for a whole cognitive economy, Combinations are mutually inclusive (any, none, or all of them may describe the set of particular cognitions that comprise a whole cognitive economy). Within a cognitive economy, Scenarios are mutually exclusive. That is, only one Scenario describes a cognitive economy, to the exclusion of all other Scenarios. It may be the case, though, that different species, or even different individual organisms within a species, may be described by different Scenarios (see 4.2.3.2).

At the end of 4.2.1, I foreshadowed the arguments for all four Combinations being instantiated in a human cognitive economy that I will prosecute in Part II. In terms

of Scenarios, this means I will argue for the Scenario that represents all four Combinations being instantiated in a human cognitive economy ($A \cup C$, below).

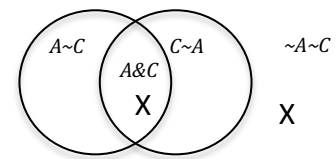
I have adopted the symbols of set theory to create a shorthand for each of the Scenarios.⁸² In the diagrams below, an X within an area of the Venn Diagram (from Figure 1 above) signifies that the Combination represented by that area of the diagram is instantiated in at least one of the cognitions that comprise that cognitive economy. The absence of an X signifies that the Combination is never instantiated in any of that cognitive economy's particular cognitions. So, for example, $A = C$ below suggests that while some cognitions involve both attention and consciousness, there are no cognitions that are attended but not conscious, and no cognitions that are conscious but not attended. An X outside both of the circles represents the idea that there are cognitions that are neither attended nor conscious, while the absence of an X outside both circles represents the idea that there are no unattended, unconscious cognitions.

This latter position is one that I have not found to be held by anyone, so while I include them in the Matrix (4.2.3.1), I have largely omitted the eight Scenarios that exclude that Combination from this discussion. Where it is necessary to refer to them, I denote them by adding an asterisk to the term for the parallel Scenario (identical, except for having an X outside the circles), e.g., $A = C^*$. For completeness, there are four other logical possibilities that few people consider likely—those with no X in the $A \& C$ intersection. Again, these are largely irrelevant to current discourse on Q, so they appear in the Matrix, but their detailed description and some discussion are in Appendix 2.

⁸² Although many authors use the language of conditional logic, I have chosen the symbols of set theory over those of conditional logic because, as should soon be apparent, it offers more tools to describe Scenarios. For example, using the symbols of conditional logic, “attention is necessary and sufficient for consciousness” (my Scenario 1), would be “ $A \Leftrightarrow C$ ”. However, there is no symbolic representation of “there is only consciousness, but never is there such a thing as attention” (STF Scenario 6).

Scenario 1. $A = C$ (Perfect Overlay)

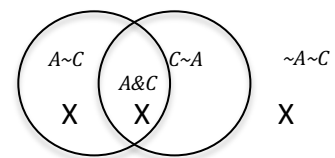
- Attention is both necessary & sufficient for consciousness. Everything that is conscious is attended (necessary). Everything that is attended is conscious (sufficient).



- Some cognitions are both attended and conscious, but there are no cognitions that are attended but not conscious, and there are no cognitions that are conscious but not attended, and of course, some are neither conscious nor attended.
- In set theory, " $P = Q$ " (the *identity condition*: every member of P is a member of Q , and vice versa) captures nicely the idea that any cognition that is attended is also conscious, and vice versa.

Scenario 2. $A \supset C$ (Attention Overflows Consciousness)

- Attention is necessary but not sufficient for consciousness (consciousness is sufficient but not necessary for attention).

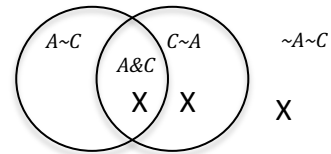


- Some cognitions are both attended and conscious; some are attended but not conscious, but all cognitions that are conscious are attended, and of course, some are neither conscious nor attended.
- In set theory, " $P \supset Q$ " (P is a *proper superset* of Q : all members of Q are also members of P , but not all members of P are also members of Q)⁸³ captures nicely the idea that all cognitions that are conscious are also attended, but not all those that are attended are conscious.

⁸³ Note that I use the symbol for a *proper* superset, " \supset ", rather than just a superset, " \supseteq ". The latter signifies that set P has more elements *or equal to* set Q , whereas here, this Scenario requires that the set of attended cognitions cannot be equal to the set of cognitions.

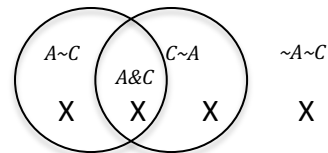
Scenario 3. $C \supset A$ (Consciousness Overflows Attention)

- Consciousness is necessary but not sufficient for attention (attention is sufficient but not necessary for consciousness).
- Some cognitions are both attended and conscious; some are conscious but not attended, but all cognitions that are attended are conscious, and of course, some are neither conscious nor attended.



Scenario 4. $A \cup C$ (Partial Intersection)

- Attention is neither necessary nor sufficient for consciousness, and consciousness is neither necessary nor sufficient for attention, but attention and consciousness sometimes occur together in the same cognition.
- Some cognitions are both conscious and attended, some are conscious but not attended, some are attended but not conscious, and of course, some are neither conscious nor attended.
- In set theory, $P \cup Q$ (the *union* of P and Q : members that belong to set P , or set Q , or to both sets) captures nicely the idea that cognitions may be either attended or conscious or both.⁸⁴



⁸⁴ More formally, $P \cup Q = \{x : x \in P \vee x \in Q\}$, where it is not true that P and Q are either identical or disjoint sets.

4.2.3.1 The Matrix

The Scenario that describes a given cognitive economy is defined by the set of Combinations that are instantiated in that cognitive economy. This relationship between Scenarios and sets of Combinations is summarised in the matrix below. Using this matrix, one can determine an author's position on Q by considering which Combinations she believes are instantiated (even rarely), and therefore, which Scenario she holds to obtain. To hold that Scenario $A = C$ obtains is to give a positive answer to Q, while to hold that $A \supset C$, $C \supset A$, or $A \cup C$ obtains is to give a negative answer to Q.

Note that this STF can be applied to the relationship between any two cognitive concepts. It can describe relationships between, say, consciousness and working memory, or attention and perception. Beyond cognitive science, it can be applied to any two entities, like gravity and mass, or anxiety and depression.

Table 3. The Matrix of all possible STF Scenarios.

<i>Scenarios</i>	<i>Combinations</i>			
	$A \sim C$	$A \& C$	$C \sim A$	$\sim A \sim C$
1. $A = C$ (<i>Perfect Overlay</i>)		X		X
2. $A \supset C$ (<i>Attention Overflows Consciousness</i>)	X	X		X

<i>Combinations</i> <i>Scenarios</i>	<i>A~C</i>	<i>A&C</i>	<i>C~A</i>	<i>~A~C</i>
3. $C \supset A$ (<i>Consciousness Overflows Attention</i>)		X	X	X
4. $A \cup C$ (<i>Partial Intersection</i>)	X	X	X	X
5. \emptyset (<i>Inattentive Unconsciousness</i>)				X
6. $A \setminus C$ (<i>Attentive Unconsciousness</i>)	X			X
7. $C \setminus A$ (<i>Conscious Inattentiveness</i>)			X	X
8. $A \Delta C$ (<i>Complete Independence</i>)	X		X	X
9. $A = C^*$		X		
10. $A \supset C^*$	X	X		
11. $C \supset A^*$		X	X	
12. $A \cup C^*$	X	X	X	
13. \emptyset^*				

<i>Combinations</i> <i>Scenarios</i>	<i>A~C</i>	<i>A&C</i>	<i>C~A</i>	<i>~A~C</i>
14. $A \setminus C^*$	X			
15. $C \setminus A^*$			X	
16. $A \Delta C^*$	X		X	

4.2.3.2 *Strong and Weak Readings*

It is also necessary to specify the domain over which a Scenario obtains. For example, it may be that among the human population, a variety of Scenarios may be found to obtain in different individuals, or perhaps vary across very different cultures. Call this the *Very Weak Reading*. Or, a Scenario may describe all normally functioning members of the species *homo sapiens*, but fail to obtain only in the most extreme cases (e.g. vegetative patients). Call this the *Weak Reading*. Or, a Scenario may obtain in all members of the species *homo sapiens* without exception, being integral to the very definition of that species. Call this the *Strong Reading*. Or, a Scenario may be the only one that obtains for any and all creatures capable of both attention and consciousness. Call this the *Very Strong Reading*. Some of the disagreements over the nature of the relationship between attention and consciousness may be clarified by considering which of these readings is being adopted by their authors—different readings may be compatible with different Scenarios, or, vice versa, the same empirical evidence may lead one to prefer different Scenarios when answering Q for the different readings.

The reading that is appropriate to the ascription of a Scenario is an empirical matter, not analytical. In this treatise I argue for a Weak Reading of the Scenario $A \cup C$ as being the one that obtains, but it is not hard to imagine a possible world in which a different reading might have been ascribed with perfectly logical coherence.

I consider the Weak Reading to be the most likely, by a process of elimination. The Very Weak Reading seems unlikely given the copious evidence that the brains of the vast majority of human beings function the same way in regard to attention and consciousness. The Strong Reading seems to be discounted by cases such as comatose patients, who, while clearly falling under a different Scenario to that obtaining in healthy human beings, are still considered to be human beings, and treated with all the respect entailed by that. For this reason, and for the additional reason that we know very little about the bizarre and wonderful cognitive economies of other species (much less those of hypothetical alien species), the Very Strong Reading seems unjustifiable. When I argue in Part II, then, that $A \cup C$ is the “right” Scenario, it is on the basis of a Weak Reading of that claim.

4.2.4 Other Taxonomies

Other frameworks that have been employed have facilitated much fruitful discussion:

“Much work has focused on the relation between consciousness and attention, usually framed in terms of the twin questions of whether attention is necessary and/or sufficient for consciousness. Every possible answer has been entertained” (Seager, 2015).

But STF provides a framework for examining Q that is complete, precise, and optimal to a greater degree than previous frameworks, none of which truly entertain “every possible answer.” By complete, I mean it encompasses all the logical possibilities for how attention might be related to consciousness. By precise, I mean that its taxonomic categories are unambiguous. And by optimal, I mean that it is the most effective way to structure the possible answers to the question. By gathering together all possible views in a single conceptual space, STF encourages comparison and contrast between them.

In this section I highlight the advantages of STF over a number of existing frameworks (Appendix 3)⁸⁵ with respect to completeness, precision, and optimality. It is important to note that most of these frameworks are not intended to be attempts at a *comprehensive* framework at all, so my analysis is not in any way a criticism.

4.2.4.1 *Complete*

A comprehensive framework needs to be *complete*. The conceptual space it maps should exhaust all logical possibilities. STF does this in two ways. First, both the Combinations and the Scenarios are exhaustive. That is, the four Combinations are the only logically possible Combinations, and the sixteen Scenarios are the only logically possible Scenarios. Second, describing both instantiations of Combinations and Scenarios that obtain given those patterns is comprehensive. That is, it covers all the territory required to answer Q across all the readings discussed above (4.2.3.2).

⁸⁵ These tables are in themselves a valuable tool for comparing the views of their authors, and the reader may find it helpful to refer to them while reading the discussion below.

What is more, a truly comprehensive framework will be one that is capable of completely subsuming other, less comprehensive frameworks (and thus providing a valuable common basis for comparison between the other frameworks and the authors' answers to Q). STF does this in both the ways mentioned above: it subsumes the partial lists of Combinations and of Scenarios found in other frameworks in its exhaustive lists; and it subsumes frameworks that describe only Combinations or only Scenarios within its two-storey structure. I consider each in turn.

Many discussions include only some Combinations without offering any sort of opinion at all about the others. For example, Hassin et al., (Hassin, 2013; Hassin, Bargh, Engell, & McCulloch, 2009) discuss $A \sim C$ at length, but make no mention at all of $C \sim A$. By comparison, Koch and Tsuchiya (2007), van Boxtel et al., (2010), and Hohwy (2012, p. 5)⁸⁶ do consider all four Combinations, thus providing a more comprehensive view of the landscape. However, their framework does not adequately address all the possible Scenarios.

Other frameworks provide a Scenario-type taxonomy, but all of these are incomplete. For example, the framework of Iwasaki (1993) seems to include just two Scenarios, $A = C$ and $C \supset A$, although a lack of precision (see below) makes it possible to also read his options as $A \supset C$ and $A \cup C$. Schwitzgebel's distinction between "rich" versus "thin" consciousness most likely includes the same two Scenarios as Iwasaki, although again a lack of precision leaves open the possibility of other interpretations. Lamme (2003) also includes just two STF Scenarios in his five categories, $A = C$ and $A \cup C$. De Brigard (2010), Montemayor and Haladjian (2015), and Pitts et al., (2018) feature just three out of the four live Scenarios, leaving out $C \supset A$.⁸⁷ Montemayor and Haladjian come closest to a truly comprehensive framework, in that one of their categories, "Full Dissociation", can be read as referring to STF

⁸⁶ On the basis of a predictive coding account of attention and consciousness.

⁸⁷ Somewhat confusingly, De Brigard enumerates only three Scenarios as being "general views" (p. 189), framed in the language of contingency. He then goes on to discuss Mole's (2008a) view, which is $C \supset A$, thus completing the set of the four live options of STF.

Scenarios 5-8, lumping the four Scenarios into just one category in which there is never any such thing as conscious attention. Their framework still falls short of a universal framework, then, since it does not distinguish between the four Scenarios where there is no conscious attention, and overall, it covers only seven out of the possible sixteen Scenarios.

This brief survey illustrates the ability of STF to comprehensively subsume existing frameworks, and gives some reason to hope that it would subsume any other frameworks not considered above, whether past or future. The value of the twelve unlikely Scenarios (i.e., Scenarios 5-16) is illustrated by having Scenarios 5-8 available to cover the odd categories such as that of Montemayor and Haladjian.

4.2.4.2 *Precise*

I noted that some of the frameworks cited above were imprecise. By *precise*, I mean that the framework is precise about the conditions for membership in, or exclusion from, its categories. The four Combinations and sixteen Scenarios of STF are quite precise in this sense. It is relatively straightforward to classify any cognition into one of the Combinations (provided one is clear about the definitions of attention and consciousness), and as the matrix above shows (4.2.3.1), relatively straightforward to derive the appropriate Scenario from any given set of Combinations.

This kind of precision would improve the clarity and utility of some of the existing frameworks. For example, because Iwasaki does not address the possibility of $A \sim C$, all his three categories are ambiguous as to which Scenario they denote. His formulation of the first category, “people may perceive only those stimuli that are under focal attention” (Iwasaki, 1993, p. 212), explicitly invokes $A \& C$, but it could also implicitly allow for $A \sim C$. This leaves it unclear whether it is $A = C$ or $A \supset C$. The second view he discusses commits to $A \& C$ and $C \sim A$, but because it doesn’t say anything about $A \sim C$ it is unclear as to whether it is $C \supset A$ or $A \cup C$. Similarly, his third view is just a variation on the second, in which attention and consciousness are independent processes that may co-occur, but without attention modulating

phenomenal character. This is compatible with the same Scenarios as his second view, but addresses the additional question of the influence of attention on conscious contents, a question I address below in 4.4. Applying STF to Iwasaki's implicit framework allows us to formulate his arguments more precisely.

Schwitzgebel's (2007) distinction between "rich" and "thin" consciousness can also be confusing. Roughly, to hold that consciousness is rich is to hold that we are conscious of more than that to which we attend, while to hold that consciousness is thin is to hold that we are only conscious of that to which we attend. But what exactly does thin consciousness mean? Applying STF to the concept suggests that it commits to the Combination $A \& C$, while rejecting $C \sim A$. But the status of $A \sim C$ is not clear. This actually leaves thin consciousness as qualifying for either the $A = C$ Scenario (if $A \sim C$ is rejected) or the $A \supset C$ Scenario (if $A \sim C$ is allowed). The tenor of the discussion in Schwitzgebel's paper suggests the former, but for the purposes of Q, we need to know for sure.

Neither is rich consciousness without ambiguity. On the face of it, it seems quite clearly to be $C \supset A$, since it allows for $A \& C$ and $C \sim A$, but not $A \sim C$ (actually, $A \sim C$ is not mentioned explicitly, but the tenor of the discussion seems to exclude it). But Cohen et al., (2016b)⁸⁸ suggest that one can embrace a rich view of consciousness that is $A = C$. They achieve this by arguing that ensembles and summary statistics are fully attended, and are sufficient to illusorily provide the richness of conscious experience (Chapter 7). Thus, consciousness and attention overlap perfectly in richness, a kind of Scenario $A = C$ because we are mistaken in thinking that there is any $C \sim A$. Here, the precision of STF has helped us to understand how Cohen et al.'s arguments relate to Schwitzgebel's by providing a common conceptual space against which they can be compared.

⁸⁸ See also the subsequent discussion generated from this paper in *Trends in Cognitive Sciences*, 2016, volume 20, number 9, pages 641-644.

STF establishes the outline of the relationship between attention and consciousness in discrete binary terms, but it can easily accommodate gradation of different kinds, and in fact, bring a higher degree of precision to their analysis. One kind of gradation is the degree to which attention (or consciousness) occur: there may be less or more attention. Thus, Koch and Tsuchiya (2007) use the term, the “near-absence” of attention. The idea of a graded attention also appears in van Boxtel et al., (2010). STF allows for gradation within its Combinations (and therefore, Scenarios). A cognition that is *A&C* may be heavily or lightly attended, however that magnitude might be measured.⁸⁹ For the purposes of Q, and as many have remarked (De Brigard & Prinz, 2010, p. 57; Koch & Tsuchiya, 2008; Mole, 2008b), the “near-absence” of attention is actually the presence of attention, not the absence of attention.

STF allows us to capture another kind of gradation—the degree of overlap between attention and consciousness. That is, where there is an overlap between attention and consciousness, the degree to which there is an overlap may be an interesting question in its own right. This can be represented graphically with the Venn diagrams of STF (Figure 2, below). This is reminiscent of the approach taken by Montemayor and Haladjian (2015) in their CAD framework, which is again comfortably subsumed by STF, and given a precise formulation.

Another kind of imprecision is inherent in the approach of De Brigard (2010) who adopts a *via negativa* of sorts. Unlike most authors, who define their views by *affirming* particular Combinations, he defines his view by *rejecting* particular Combinations. STF can accommodate this approach too. We can think of each of De Brigard’s rejections as deleting an X from the Venn diagram (or from the Matrix). The cause of the ambiguity inherent in De Brigard’s approach now becomes obvious. By remaining silent on the Xs he neither affirms nor denies, he leaves open multiple possibilities for Scenarios. Of course, he deals with this ambiguity in his paper in

⁸⁹ For example, if one defines attention as being the devotion of cognitive resources above a certain threshold, one can determine a value for an attended cognitive process that reflects quantitatively how much above that threshold it sits.

other ways and makes his preferred Scenario quite clear (it is $A = C$), but the danger of imprecision inherent in a *via negativa* is nicely evoked by applying STF.

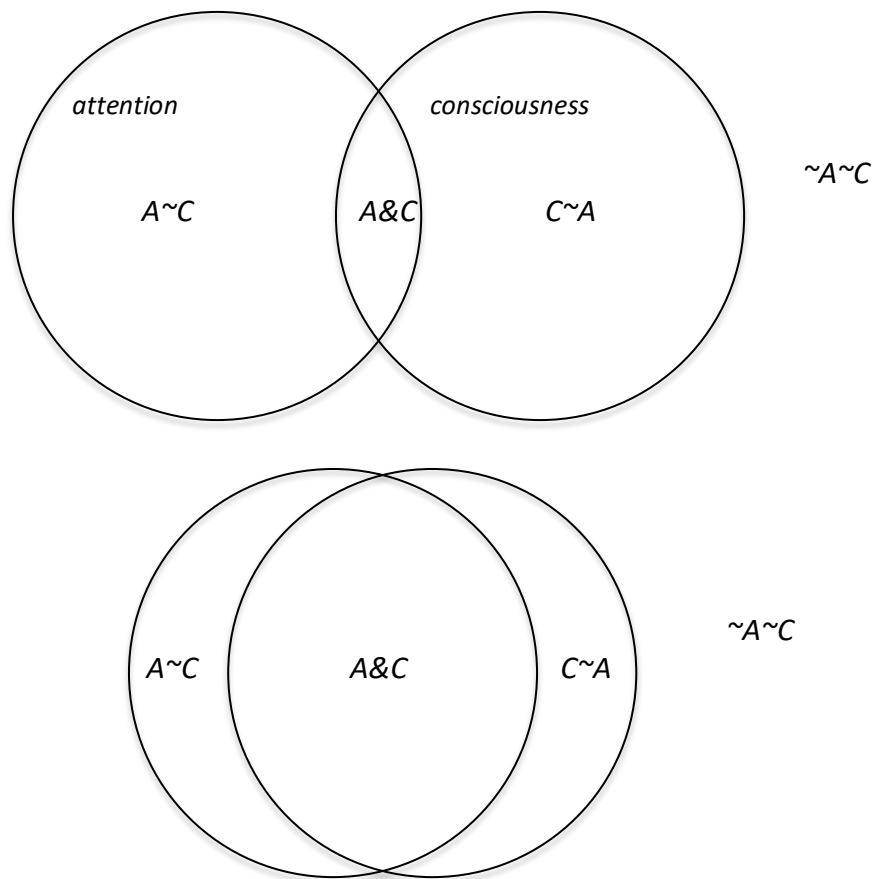


Figure 2 A graphical representation of one way that gradations in the overlap between attention and consciousness might be represented. *Above*, very little overlap; *below*, a great deal of overlap.

It could be argued that the precision of STF is not really desirable, that STF lacks a certain nuance. For example, we have seen that some approaches focus on determining the relative *influence* of attention on consciousness (e.g. Iwasaki) or on the *degree* of overlap between attention and consciousness (e.g. Montemayor and Haladjian). These kinds of nuanced relationships are not captured intrinsically by

just considering Combinations and Scenarios. Does this count against the usefulness of STF? I would argue that on the contrary, I have shown that STF acts as a basic foundation upon which other approaches, like those just mentioned, can be built. What is more, STF acts as a single conceptual space within which different, more nuanced approaches can all be located and compared. This common space makes it easier to understand how those different approaches relate to each other.

Another possible criticism of STF is that its precision is too restrictive. To force all views on the topic into one or other of its Scenarios could result in a stifling of imagination and creativity. What if there are useful views that just don't fit into STF at all? In response, I point out the ability of STF to subsume quite a wide variety of imaginative and creative approaches to the topic. What is pertinent is that when they are fitted in to STF, their originality is elucidated rather than attenuated. STF is a precise analytic framework that captures the essential nature of the conceptual space effectively, allowing quite varied accounts to be built upon it.

4.2.4.3 *Optimal*

Employing so complete and precise a framework as STF confers a number of benefits in the quest to answer Q—it is, of the available frameworks, the *optimal* framework. Many of the authors mentioned above discuss only one or more Combinations without touching upon Scenarios. Of course, they are not all attempting to answer Q. But this makes it clear that given the broad affirmation of the Combinations $A \& C$ and $\sim A \sim C$, Q can only be answered when the status of remaining two Combinations— $A \supset C$ and $C \supset A$ —have been determined within a cognitive economy. Anything less than that leaves Q an open question. STF clearly maps out the path that must be taken if we want to see the wood, not just the trees.

STF is also the optimal framework in that it provides the common ground upon which the views of different authors, often employing quite different language, can be standardised and compared. It has the potential to become a *lingua franca* that unites disparate approaches. It allows us to see, for example, that while Montemayor

and Haladjian's (2015) consideration of Combinations is not as well-structured as that of Koch and Tsuchiya (2007), they do consider a broader range of possible Scenarios than De Brigard (2010). But STF still has the advantage over Montemayor and Haladjian's CAD. It not only provides the apparatus to quickly and easily situate it against the other frameworks, it also subsumes its seven Scenarios within the sixteen of STF. Interestingly, the only Scenario Montemayor and Haladjian omit from STF's first eight is $C \supset A$. They themselves admit that their "framework is not entirely neutral" (p. 21), being influenced by their assumptions about neural correlates and evolutionary considerations. STF, I would argue, is based on no such pre-assumptions and ought to be useful for projects such as comparing competing models and developing new ones, independently of any theoretical assumptions. Further, while they discuss the possibility of $C \supset A$, they leave it out of their framework on the basis that the $C \sim A$ Combination is a difficult thing to prove empirically. This difficulty is a significant one, and indeed, plays a large role in the discussions of Part II of this treatise, but a comprehensive analytic framework must include all hypothetical possibilities and leave their confirmation to empirical investigation.

4.3 What Others Have Said

Having established STF as a framework for addressing Q, I turn in this section to reviewing some of the vast literature relating to Q, situating various authors against STF, and sketching some pertinent observations on the discourse. In Appendix 4 I classify authors by allocating them to one of the four live Scenarios with just a brief comment on each.⁹⁰ While I have done my best to identify the Scenarios of the

⁹⁰ Phillips (2011a, p. 222 endnotes 6 & 7) has a much briefer survey.

authors as accurately as possible, this is often a difficult thing to do. As observed above (4.2.4.2) some authors do not address enough Combinations to confidently identify their preferred Scenario. They may also mean different things by the terms “attention” and “consciousness,” or use different terms altogether (e.g., “access” or “awareness”), which may partially explain why they prefer different Scenarios. And, I have insufficient space to fully present the nuances of their arguments for their preferred position, and must mostly be content with categorisation without much explicit justification.

In this section I selectively reflect on some of the work of the authors surveyed in Appendix 4 in light of the STF, drawing observations pertinent to my arguments in Part II.

4.3.1 STF Classification

I have argued that STF brings a degree of clarity and precision to the views of authors in the discourse about Q. Some authors are easy to allocate to a Scenario, but others are less so. Of course, many authors considered here are not specifically addressing Q, but issues that are related to Q in some way, so their ambiguity in relation to Q is not in any way a criticism. But for the researcher interested in Q, applying the STF can be helpful, and even promote a deeper understanding of the authors’ account. I consider some examples of authors that are easily classified under STF, and then a more ambiguous example.

Crick et al., (Crick & Koch, 1990; Koch & Tsuchiya, 2007; Tsuchiya & van Boxtel, 2010; van Boxtel et al., 2010) present an unambiguous case for $A \cup C$. In addition to the four KT arguments discussed in Chapter 1, the idea of *fleeting awareness* (Crick & Koch, 1990, p. 272)— a neurally and functionally distinct fast process linked to iconic memory, as opposed to the slower process of binding neurons via visual attention to produce a single visual object to short-term memory—provides a neural

underpinning to $C \sim A$.⁹¹ They also hold that while attention is the mechanism that binds content into working memory, “the attentional mechanisms themselves are largely unconscious” (p. 269), which is $A \sim C$. The $A \cup C$ Scenario is also unambiguously championed by Block, Kentridge et al., and others.

The $A \supset C$ Scenario finds unambiguous champions in De Brigard and Prinz and O'Regan and Noë, but the only unambiguous proponent of $C \supset A$ I could find was Smithies. Smithies' argument is that it is a functional characteristic of attention that it “makes information accessible for use in the rational control of thought and action ... there is no attention without consciousness, since no unconscious information is fully accessible for use in the rational control of thought and action” (Smithies, 2011, p. 248). However, this definition of attention seems to be inadequate for his argument. Not being “fully accessible” is not the same thing as not being accessible at all.⁹² If even some information can be made available to the processes of rational control without ever being conscious, that would be enough to establish $A \sim C$. And that would seem to be the case in, for example, priming, where masked or otherwise unconscious information does in fact measurably influence behaviour or report (6.2).

Other authors are much more ambiguous, and require some detective work.⁹³ Posner's (1994) creative approach draws an analogy between the relationship between attention and consciousness and the relationship between DNA and “life.” While DNA is not itself “life,” one cannot fully understand life without understanding DNA. He suggests that attention bears the kind of constitutive or causal relationship to consciousness that DNA bears to “life” (I consider these kinds of relationship below in 4.4). However, it seems implausible to suggest that DNA is necessary for life (if strong Artificial Intelligence is correct for example) or that it is sufficient for

⁹¹ See also 8.5.

⁹² Compare this response to De Brigard & Prinz's response to Koch & Tsuchiya (2007): “the near absence of attention is not the same as the absence of attention” (De Brigard & Prinz, 2010, p. 57).

⁹³ It is noticeable that these tend to be older texts. Perhaps later authors, responding to the more explicit posing of Q over time, took pains to be clearer on their position on Q.

life (undenatured DNA in a recently deceased corpse for example). Therefore, while some (Cohen, Cavanagh, Chun, & Nakayama, 2012, p. 411; Hardcastle, 1997, p. 59) have attributed an explicitly $A = C$ position to him, I prefer to be a little more cautious. Posner's discussion of how the putative neural correlates of attention might fit into the contemporary models of consciousness of Crick, Edelman, and Beck & Eccles (Posner, 1994, p. 7403) is further reason for considering him implicitly $A = C$ rather than explicitly.

Mack and Rock (1998) pioneered the study of *inattention blindness*—the phenomenon where, when one's attention is focused on a particular object or task, quite obvious stimuli peripheral to that task completely evade conscious experience.⁹⁴ Ahmadi et al., (2011, p. 1366) take them to reject the possibility of $C \sim A$, on the strength of statements such as this: “there is no *conscious* perception without attention” (Mack & Rock, 1998, p. 14, italics in the original). Concurring, De Brigard & Prinz (2010, p. 52) comment that, for Mack & Rock, “every occurrence of consciousness is going to be, as a matter of definition, an occurrence of attention.” This rejection of $C \sim A$ leaves open the question of which of the Scenarios $A = C$ or $A \supset C$ Mack and Rock would be committed to. Cohen et al., (2012, p. 411) class them among those who hold “that attention and awareness are inextricably linked”, which sounds very much like $A = C$, and Schwitzgebel (2007, p. 7) takes them to espouse the thin view, which is vaguely $A = C$.

However, on the question of $A \sim C$, Mole (2008a, p. 93) says that “Mack and Rock claim ‘that attention is necessary for conscious perception’ (p. 250)”, but that they “reject the claim that consciousness is necessary for attention” (p. 93). In support, he cites them thus:

“Unfortunately, although the proposal that conscious perception and attention refer to identical processes has the advantage of

⁹⁴ E.g., the invisible gorilla in plain sight of Simons and Chabris (1999).

simplicity, it is discredited on several grounds. First, it would appear to lead to the false conclusion that there can be no attention without [conscious] perception. This conclusion seems false on both experiential and empirical grounds” (Mack & Rock, 1998, p. 245).⁹⁵

Thus, Mole argues, they unambiguously embrace $A \sim C$, (and for my purposes, are therefore $A \supset C$). However, I argue this is too strong a conclusion to draw, and they may not in fact be embracing $A \sim C$. Consider two of their four reasons for rejecting the equation of attention with perception—anticipatory attention and vigilance—in light of the Target Question (1.4.6). In these cases, attention Targets not the *object* that is ultimately identified when it appears, but the *empty perceptual field* into which the object eventually enters. The subject is both conscious of and attending-as-vigilance to the field itself. And before the object appears, she is both unconscious of and not attending-as-vigilance to the object itself. There is, in fact, no $A \sim C$ here. Mole himself discusses the situation where one attends to a blank radar screen in anticipation of a blip appearing. He rightly points out that this is not a case of attention without conscious perception, for it is “*by perception* that the vigilant radar operator knows that no pip has occurred ... it confuses the perception of absence with the absence of perception” (p. 98). Another way to put this is that the operator exercises *spatial* attention in the absence of *object* attention (since there is no object to attend to). Thus, Mack and Rock’s statement above suggests that their position on $A \sim C$ may be open to revision, and this ambiguity justifies our keeping them tentatively under the $A = C$ Scenario.

⁹⁵ The “[conscious]”, inside square brackets, is inserted by Mole when he quotes them, but not present in their original. It is a plausible insertion, however, since they use the term “conscious perception” at the beginning of the quote.

4.3.2 Using Q to Produce Cognitive Models

Q not only an interesting question in itself, but has also contributed to our understanding on other questions. Although Dehaene et al., (2006) and Lamme (2003) subscribe to different Scenarios, they both build aspects of their cognitive models at least partially on views acquired from considering Q. Dehaene et al.'s threefold distinction between *subliminal* (never conscious); *pre-conscious* (potentially but not actually conscious); and conscious processing arises out of their grappling with how attention and consciousness interact, as do Lamme's four models of the temporal progression of processing (p. 13, Figure 2).

Emphasising the importance of the temporality of attention and consciousness, Wolfe's (1999b) exploration of the interplay between attention and consciousness in the temporal pattern of visual processing in vision underwrites his distinction between *pre-attentive* and *post-attentive* consciousness. He argues that the visual content outside the focus of selective and conscious visual attention—the background—is itself conscious, but rather data-poor. This is pre-attentive consciousness. Should the focus of attention turn to that region of the visual field, the content is stabilised, integrated, and interpreted, and becomes post-attentively conscious, while the region that was previously post-attentively conscious now falls back into pre-attentive consciousness once more. I consider this kind of temporal interplay in more detail in Chapters 7 and 8, where I propose my own model of cognitive processing.

4.3.3 Answering Q by an Act of Definition

The answer one derives to Q depends heavily on how one defines attention and consciousness. In Chapter 3, I argued against Phenomenal Definitions of attention as

begging the question of Q. There may be more subtle examples of this problem. For example, O'Regan and Noë (2001)—who Cohen et al., (2012, p. 411) read as $A = C$ with good reason—state at one point: “But if you should turn your attention to the color of the car ahead of you, and think about it, or discuss it with your friend, or use the knowledge of the car’s color to influence decisions you are making, then, we would say, you are aware of it” (p. 944). The first part of that statement, before the word “then,” describes attentional processes. They then assert that the presence of any of these manifestations of attention justifies us imputing the concurrent presence of awareness.⁹⁶ This assumption of a Phenomenal Definition of attention is based on the clear case of $A \& C$ they describe, but makes no allowance for other cases, such as the $A \sim C$ of being unaware that you have lifted your foot of the accelerator as you passed a high-visibility speed camera while engaged in deep conversation with your passenger.

More subtly, it could be argued that Block (1995) might be charged with locking himself in to $A \cup C$ the moment he commits to his distinction between access consciousness and phenomenal consciousness. By drawing this conceptual distinction he thereby gives rise to the possibility—for which he argues—of phenomenal content that is not accessed, and accessed content that is not phenomenal. But Block takes these questions as necessarily empirical ones and answers them accordingly, which contrasts with the answer-by-definition strategy of O'Regan and Noë.

Mogensen and Overgaard (2018) derive different answers to Q depending on definitions. Based on their *reorganization of elementary functions* framework, they conclude that while phenomenal content can overflow working memory (which is closely tied to attention and access, see Chapters 5 and 7) it cannot overflow the availability of information for action (which is also a possible definition of attention).

⁹⁶ A synonym for phenomenal consciousness in this context.

4.3.4 Influence of a Prior Theoretical Basis

This highlights another point, that the theoretical framework from which one works also strongly influences one's answer to Q. If one thinks, as Lamme (2003, 2010) does, that the neural correlate of consciousness is *recurrent neural processing* (or RP—roughly, waves of both feed-forward and feedback activation in the brain), and that such RP occurring only in the posterior brain is sufficient for consciousness, then consciousness may arise in the absence of the attention that is generally thought to depend upon involvement of the frontoparietal regions. This is $C \sim A$. And since he further considers attentional modulation of the *feed forward sweep* in the absence of feedback activation to be possible (Lamme, 2003, p. 16), he must allow for $A \sim C$. Lamme's recurrence theory of consciousness underwrites his preference for $A \cup C$.

Prinz's *Attended Intermediate-level Representation* theory of consciousness (Prinz, 2011, p. 182) leads inevitably to $A = C$. Attention is the gateway to working memory, modulating intermediate-level representations,⁹⁷ and making them available to working memory. All and only content that is processed in working memory in the right ways becomes the content of consciousness—hence the necessity and sufficiency of attention for consciousness, via working memory.

⁹⁷ An idea Prinz builds on earlier models, e.g., Jackendoff (1987).

4.3.5 Interpretation of Empirical Work

In recent decades, it is rare to find literature addressing Q that does not engage deeply with empirical research. But proponents of different views often differ as to how to interpret that research. In 1.4 I discussed some aspects of the empirical syntheses of KT and DP. Here I look at another central strand in the discourse about Q, the question of the relative *grains* of attention and consciousness.

I have classified Block as espousing $A \cup C$. While he mainly focuses on $C \sim A$, based largely on the foundational work of Kentridge et al., (see Appendix 4) and others,⁹⁸ he has also argued clearly for a double dissociation between attention and consciousness (Block, 2013b, p. 182; Block et al., 2014, p. 556), thus committing also to $A \sim C$. Much of Block's recent work on $C \sim A$ deals with phenomenal overflow—the idea that phenomenal consciousness overflows access consciousness, a topic I address in detail in Part II. In an interesting variation on this theme, Block (2013b) has argued that attention and consciousness seem to be of different “grains”—different resolution or richness of detail. That is, there is empirical evidence to suggest that the content of phenomenal conscious is of a much finer grain than is the content of attention.⁹⁹ This has provoked some fascinating discussion, often specifically about the different interpretations of the empirical data that are possible (Block, 2013a, 2014b, Richards, 2013, 2015, 2016; J. H. Taylor, 2013b; Tye, 2014), some aspects of which I return to later.

The debate raised by Block's grain argument illustrates the remarkable propensity to interpret the same empirical work in different ways to come to contradictory conclusions. This phenomenon is also evident in the debate inspired by KT's empirical case for $A \cup C$ and the challenging of their interpretations by DP, whose interpretation is in turn challenged by Barrett (2014). There are, of course, many complex reasons for this divergence of interpretation, and much work has been

⁹⁸ For a challenge to the interpretations of empirical evidence advanced in support of overflow by Block and others, see Gross and Flombaum (2017).

⁹⁹ Or access consciousness (7.2.2).

done in the field of scientific objectivity that shows that investigators are rarely ideal objective interpreters (Reiss & Sprenger, 2017). I discussed the importance of structuring the logic of one's interpretations carefully in 1.4.7. Another interesting question is whether an investigator's field might affect their preferred Scenario.

4.3.6 Influence of Background

The literature touching on Q is vast. While my survey (Appendix 4) of nearly fifty authors or groups of authors is in no way exhaustive or rigorous, and no strong conclusions can be drawn from it, it does reveal some interesting patterns.¹⁰⁰ I summarise a simple statistical survey of authors by Scenario in Figure 3 below.¹⁰¹ While there is no clear “majority” view, $A = C$ and $A \cup C$ are significantly more popular than the other two live Scenarios. When we come to consider the academic fields of authors, we find that while both these Scenarios have a roughly equal number of proponents among the philosophical and non-philosophical camps, $A \supset C$ and $C \supset A$ show a more skewed pattern. No recent non-philosophers have espoused $C \supset A$ (both Wolfe and Iwasaki were writing in the 1990s), and only one philosopher (Richards) has espoused $A \supset C$.

¹⁰⁰ One startlingly obvious pattern is that while there is a spread of authors geographically and even across disciplines (including medicine, social science, mathematics, ophthalmology, etc., see Appendix 4, Table 14), there is a remarkable gender imbalance among lead authors. Of the 48 here classified, only 4 are female (92% male): Arien Mack ($A = C$); Valerie Hardcastle ($A \cup C$); Catherine Tallon-Baudry ($A \cup C$); and Carolyn Dicey Jennings ($C \supset A$). This is all the more remarkable given that the vast majority of discourse on Q surveyed here occurred in the last 30 years, rather than being spread over centuries, and is likely an extreme manifestation of the gender imbalance in philosophy more generally (Holtzman, 2016; Hutchison & Jenkins, 2013) and in neuroscience (González-Álvarez & Cervera-Crespo, 2017).

¹⁰¹ This, by the way, illustrates another benefit to STF. It makes this kind of meta-analysis much easier.

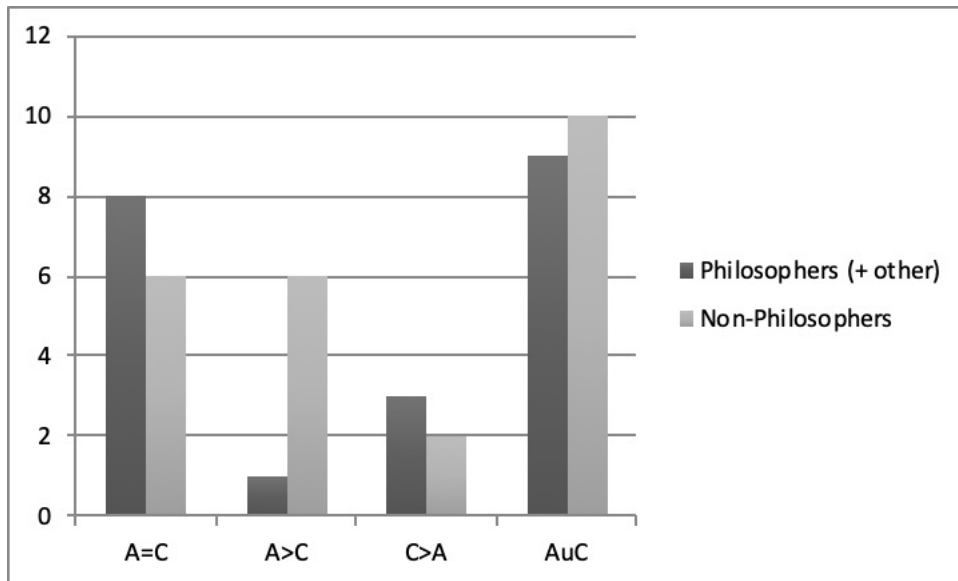


Figure 3 The preponderance of philosophers vs non-philosophers among the Scenarios of STF. The data and some notes on methodology for this survey may be found in Appendix 4.

The figures are not statistically significant, but there may be good reason behind this pattern. It almost certainly lies in the history of the study of the mind over the past hundred years or so.

“Most psychologists working on attention, because of the intractability of the problem of consciousness, had either no interest in consciousness or no way to connect their findings with considerations about consciousness.” (Montemayor & Haladjian, 2015, p. 18)

“Studies of attention are a mainstay of contemporary cognitive science. Understanding the mechanisms of attention has proved to be essential for understanding a range of information-processes, not least the processes of vision. By comparison, and

notwithstanding some recent work, attention has been a peripheral concern for philosophers” (Stazicker, 2011b, p. 1).

One plausible story might go something like this. Philosophers have spent a lot of their time over the years pondering problems to do with phenomenal consciousness, and very little time pondering attention. When they finally turned their minds to the question of the relationship between the two, some of them at least were somehow “primed” to see consciousness everywhere, but less likely to recognise attention. Thus, in the event they rejected $A = C$ and $A \cup C$, they were thus far more likely to end up espousing the Scenario where consciousness predominates, namely $C \supset A$. A similar story may be said of non-philosophers: psychologists, neuroscientists, biologists, and the like, in which fields speaking of “consciousness” was for some time almost a taboo. No doubt, other stories may be told. This seems like a fascinating area for future research. As I have admitted, the small sample size of this survey limits its statistical significance, but it would be interesting to have a question like this included in any future large scale surveys of the views of philosophers, like those of Bourget and Chalmers (2014), or that conducted by *Philosophy Now* (2012).

4.3.7 Consensus Position?

The spread of authors across the Scenarios justifies the inference that Q remains an open and highly controversial question. But if there is no *consensus* on an answer to Q, can we at least identify a *predominant* position? This too turns out to be a vexed issue. Is a position predominant among “the folk,” or among scientists generally, or among philosophers of mind, or among that small group of people who have researched the specific question of the relationship of attention to consciousness? We must also beware lest we fall prey to the fallacy of *argumentum ad populum*—majorities often get things wrong. The search for a predominant position can be

complicated by the desire of some authors to overemphasise the position they are arguing against in order to make their arguments braver and more daring. With these and other caveats in mind, it may seem that the search for a predominant position is something of a fool's errand. But we can glean some useful insights from the attempt. There are many opinions on what constitutes the predominant position on Q in the literature. For example,

“psychologists have been increasingly confident that it will turn out that everything that one is conscious of is a thing to which one is attending” (Mole, 2008a, p. 93).

“These results led many to suggest that attention and awareness are inextricably linked. Broadly speaking, this was a standard assumption for many years” (Cohen et al., 2012, p. 411).

“Most researchers closely link attention with awareness (equated here with the contents of conscious experience), arguing that the two always occur together. That is, attending to an object is the same as becoming conscious of it” (Koch & Tsuchiya, 2012, pp. 103–104).

This suggests that authors from a variety of Scenarios and fields see that the predominant view, the incumbent view only recently being challenged, has been $A = C$ (at least for the past thirty years). Another example of this is Allport, who goes as far as to suggest that in the world of psychology at least, attention was thought of almost as being identical with consciousness.

“The study of ‘attention’ began in phenomenology. ‘Everyone knows what attention is’, wrote William James in 1890. Ninety years later the word is still used, by otherwise hard-nosed

information processing psychologists, as a code name for consciousness (Allport, 1980, p. 113)."

But on the other hand others have portrayed the debate as being more equally balanced (De Brigard, 2010, p. 190; Schwitzgebel, 2007, p. 7). These authors have highlighted the continuing disagreement that characterises the contemporary discourse.

"It has been repeatedly maintained that attention facilitates conscious perception. However, it remains controversial whether or not the region outside of the attention spotlight is consciously perceived." (Iwasaki, 1993, p. 229)

"Yet, a minority tradition in psychology, going back to the 19th century, emphasizes that attention and consciousness are related yet different, and that one can attend to an object or feature of an object without becoming aware of it." (Koch & Tsuchiya, 2012, p. 104).

And then there are those, like Mole (2008a, p. 86), who presents the "commonsense view"—contrary to the view of the "psychologists"—as being $C \supset A$: "According to commonsense psychology, one is conscious of everything that one pays attention to, but one does not pay attention to all the things that one is conscious of." De Brigard (2010) challenges Mole's characterisation of the "commonsense view." And there are those who consider $A \supset C$ to be the predominant view:

"There is a near consensus on the claim that attention is necessary for consciousness ... [but] there are many who believe attention is

not sufficient for consciousness—that attention can be deployed for mental states which are never conscious” (Barrett, 2014, p. 9).

And then there are those who consider the predominant view to be $A \cup C$.

“Investigations into the relationship between attention and awareness appear to agree on one thing; the former is neither necessary nor sufficient for the latter” (R. Hine, 2015, p. 52).

Or that the case has been comprehensively made for $A \cup C$.

“We now know, contrary to many people’s introspective intuitions, that attention and awareness are dissociable: attention of various types can function in the absence of consciousness and there is some evidence that there may be conscious experience without attention or report. We now have an idea of the kinds of cognitive and perceptual processing that can occur in the absence of awareness, and how these may differ from conscious processing” (Block et al., 2014, p. 556).¹⁰²

Finally, there are those who take a somewhat more pessimistic view.

¹⁰² The opinion piece from which this quote comes has attached to it the names of no less than eight distinguished researchers in the field, many of whom I have classified above as proponents of $A \cup C$.

“‘Everyone knows what attention is’ is one of the most oft-quoted remarks ever made by William James (1890). As times goes by, researchers are becoming more and more sceptical of the veracity of this claim, and some have even argued that ‘it would be closer to the truth to say that “nobody knows what attention is”’ (Pashler, 1998; Styles, 1997, p. 1). I think philosophers of mind could benefit from this scepticism. The reviewed studies suggest that there is little agreement, in so far as our commonsense psychology is concerned, about the relation between attention and consciousness” (De Brigard, 2010, pp. 199–200).

Perhaps what we see here is a pattern of growing disagreement over time. What was once taken predominantly, especially in psychology circles, to be the received view, $A = C$, has now come into question so much that we can no longer speak sensibly of a “predominant view.” 58% (7 out of 12) of the authors or groups of authors cited in favour of $A = C$ published their views before the year 2000, compared to only 28% (5 out of 18) of those in favour of $A \cup C$. If anything, the tide appears to be gently turning towards a growing popularity of $A \cup C$.

4.3.8 Usefulness of STF

We have now seen the STF in action. It is worth recapping the benefits it brings to the study of the relationship between attention and consciousness, benefits illustrated in the foregoing discussion. STF has provided an exhaustive and precisely partitioned conceptual space against which the views of every author considered have been situated and can be more easily and fruitfully compared. Its application helps to clarify the initially ambiguous views of some authors on Q , and provide a better understanding of what they entail.

While almost all the views in the literature fall under one of the four chief Scenarios, the availability of the other Scenarios has also proved useful in the evaluation of views at the fringes of the debate. Further, discussion of Scenarios, rather than just some Combinations, has helped put specific empirical findings and particular individual arguments into context, and helped to reveal their significance to Q. Finally, this survey structured upon the STF has brought into sharp relief some interesting overall patterns in the positions of authors, patterns worthy of further investigation for the light they may shed on the meta-topic of how or why an investigator comes to prefer a particular position over others.

4.4 Kinds of Relationship

The STF developed above merely describes the possible *patterns* of relative occurrence of two entities, as I said earlier. What it does not capture are the *reasons* for these patterns. In this section, I survey the logically possible reasons for the patterns, providing options from which we can choose to describe more fully the relationship between attention and consciousness. My approach is relatively simple, but sufficient for the purposes of Q, without delving into relationships at different levels (e.g., Marr, 1982), or deeper philosophical theorising about the nature of causation (e.g., Woodward, 2003, 2014).

4.4.1 The Options

There appear to be four possible types of relationship between two entities such as attention and consciousness.¹⁰³ In this section, I outline and illustrate each of them, mainly to provide options from which to choose based on the arguments presented in Part II, although I will make some initial tentative conclusions on the topic here. The four kinds of relationship between attention (*A*) and consciousness (*C*) are:

Table 4. Four kinds of relationship between attention and consciousness.

- **Identity Hypothesis (IH)**
 - *A* and *C* are two names for the same thing.
 - *C* fully constitutes *C* and *C* fully constitutes *A*.

- **Partial Constitution (PC)** (or partial identity)
 - *A* and *C* share parts, but each (or at least one) has at least one part that is not the other.
 - Subtypes:
 - *A* partially constitutes *C*.
 - *C* partially constitutes *A*.
 - *A* and *C* partially constitute each other.

¹⁰³ Compare Kozuch and Kriegal's (2015) similar discussion of possible reasons underlying the neural correlates of consciousness, and Ylikoski's (2013) discussion of causation and constitution.

- **Causation (CA)**

- *A* and *C* are connected by a causal chain of some type.¹⁰⁴
- Subtypes:
 - *A* causes *C* ($A \rightarrow C$).¹⁰⁵
 - *C* causes *A* ($C \rightarrow A$).
 - Both $A \rightarrow C$ and $C \rightarrow A$.
 - Complex causal chains: e.g., *X* causes *A* and *C*; or *X* causes both *Y* and *Z*, *Y* causes *A*, *Z* causes *C*; etc.

- **Mere Correlation (MC)**

- *A* and *C* are correlated, co-instantiated, or co-occur without identity, constitution, or causation.

4.4.2 Which Option?

Which of these options is the one that describes the overall relationship between attention and consciousness within a cognitive economy? How do we even go about answering that question? Some of these options can be supported or discounted on purely analytical bases. But mostly, choosing between these options is an empirical matter, requiring empirical evidence.

An important point to note is that some of these four options exclude the possibility of others. If MC obtains, then—by its very definition—none of the other three can possibly obtain. If IH obtains, then PC and MC cannot obtain. Some form of CA might

¹⁰⁴ Block (2007, p. 482) also distinguishes causation from constitution.

¹⁰⁵ The arrows signify the concept of causation generally without engaging in the deep waters of its metaphysics.

obtain, depending on how we understand causation and constitution. Similarly, if PC obtains, then IH and MC cannot obtain, but some form of CA might obtain, depending on how we understand causation and constitution. And if some type of CA obtains, MC cannot obtain, but IH or some form of PC might obtain, depending on how we understand causation and constitution. Keep in mind that although I am addressing the relationship within a whole cognitive economy, there is no *a priori* reason why there should not be different relationships in different regions of that economy, and there may not be a single description that captures the whole cognitive economy accurately.

Whilst I aim to arrive at a reasonably confident answer as to which Scenario obtains for attention and consciousness, I have far less confidence about arriving at a confident answer as to the kind of relationship—the *reason* in some sense, behind this Scenario—in this treatise. We know too little at this stage about the brain's inner workings. The neural correlates of both attention and consciousness are yet to be definitively identified. The complexity of neural interactions—like biological systems more generally—defies lawlike description (Mitchell, 2000; Smart, 1959). And there remains a yawning explanatory gulf between mechanical descriptions of brain activity and psychological phenomena—at best, we have at present correlations, but little in the way of genuinely explanatory reductions.

4.4.3 Relevance to Q.

I content myself, therefore, with some simple observations as to how the options above relate to Q, and the employment of these observations in the rest of this treatise where they are helpful.

IH makes Q trivially true—if attention and consciousness are just the same thing, then of course, each is necessary and sufficient for the other. We saw above that while $A = C$ is a viable Scenario, it is by no means the only one, and even those who

espouse it do not, by and large, think of attention and consciousness as being *identical*. For example, *A = C* Posner (1994) explains the relationship of attention and consciousness using the analogy of the relationship between DNA and life (4.3.1), which are certainly not identical, even if in our earth-bound experience, they universally co-occur.¹⁰⁶

MC makes Q trivially false—if there is no constitutive or causal connection between attention and consciousness (and there is no other kind of substantive connection than these two), then whenever they correlate, they do so coincidentally, totally contingently. There is no necessity or sufficiency binding them together. Such a situation would indeed be remarkable, given the frequency with which attention and consciousness co-occur, so I take this option also to be highly implausible.

PC and CA bear in much more complex ways upon Q, more complex than I can fully discuss here. For example, if attention partially constitutes consciousness, then it may be that consciousness—or some “part” thereof not constituted by attention—can occur without attention. But it may also be that, given the specific features of human cognition, this never actually happens, and every instance of consciousness does indeed co-occur with an instance of attention. Life requires much more than just DNA, but it may be that without DNA, life never in fact occurs (at least on earth). And the complexity of the infinite number of possible patterns of causal chains under CA makes any of the Scenarios possible under this option.

¹⁰⁶ There are of course life forms that depend on RNA instead of DNA, but assume for the purposes of this discussion that “DNA” covers any kind of genetic nucleic acid.

4.5 Chapter Summary

In this chapter, I developed a STF for classifying the possible patterns of co-occurrence of attention and consciousness. Any particular cognition may be characterised by one of four Combinations, and within a whole cognitive economy, these four Combinations may be combined in any one of sixteen possible Scenarios, of which, four are live options as answers to Q. I argued that STF is a more complete, precise, and optimal framework for addressing Q than any existing framework, and applied it to review the literature on Q, bringing (I hope) a greater degree of clarity and facilitating comparison between the views of different authors. To better elaborate the nature of the relationship between attention and consciousness—the reason for an answer to Q—I described four broad kinds of relationship and their relevance to Q.

Employed together with the definitions of consciousness (Chapter 2) and attention (Chapter 3), STF and the four kinds of relationship (this chapter), provide the framework for my pursuit of an answer to Q in Part II. But there remains one more foundation to lay first. In the next chapter, I recruit the oft-invoked concept of working memory.

5 Working Memory

5.1 Introducing Working Memory (WM)

“Selective attention and working memory (WM) have traditionally been viewed as distinct cognitive domains ... However, a growing number of psychological and neuroscientific studies have revealed extensive overlap between these two constructs” (Gazzaley & Nobre, 2012, p. 129).

“The proposed link between working memory and conscious awareness also represents a lively and exciting interface” (Baddeley, 2003, p. 837).

“The consensus among WM- and consciousness researchers is that the psychological processes that underlie WM, attention and

awareness are closely intertwined ... The nature of these relations, however, is less consensual” (Hassin et al., 2009, p. 666).

As the quotations above show, the concept of working memory (WM) is intimately associated with both attention and consciousness and pops up regularly in Q-related discussions. In this chapter, I briefly describe the rise of WM models and metaphors, describe them, and probe them a little to extract from them some highly relevant observations and pertinent Questions that will, to differing degrees, contribute significantly to my arguments in Part II. These Questions are by no means original, but they serve to focus my thinking in the arguments that follow. I will only be able to address one of them in substantial detail in Chapter 8—the Capacity Question (5.3.2.3). Later, I employ WM to—among other things—spell out the temporal nature of both consciousness and attention, provide a principled way for dissecting out Executive Attention from Liberal Attention, and pave the way for defining the capacity limitations of Executive Attention implemented in the executive of WM.

5.2 History

“We may someday hope for a unified science of memory, but that day is not yet at hand ... the great truth of the first 120 years of the empirical study of human memory is captured in the phrase ‘it depends.’” (H. L. Roediger, 2008, pp. 227, 228).

A full understanding of WM situates it against the concept of memory generally. The nature of memory generally has baffled humans at least since the Greek poet Hesiod.

Plato's metaphor of memory as an impression made in soft wax,¹⁰⁷ described as "probably the most influential image in the entire history of discourse about memory" (Danziger, 2008, p. 28) was the forerunner to many other metaphors through history, some of which I discuss below, continuing a long and venerable tradition. Historically, memory was also a topic of major interest to Aristotle, Augustine, Avicenna, Averroes, Albert the Great, Aquinas, and also to philosophers whose initial is not "A:" Descartes, Hobbes, Locke, Hegel, Schelling, Fichte, Nietzsche, Bergson, Husserl, Heidegger, Deleuze, and Derrida (Nikulin, 2015).

In modern times, defining memory seems to have become more complicated, rather than less. Roediger (2008) points out that the optimistically simple laws of learning and memory proposed by researchers in the early twentieth century dissolved away in the face of increasingly sophisticated experimentation and theoretical work, so that he is able to state that "no general laws of memory exist. All statements about memory must be qualified" (p.227). Even a "law" as intuitive and basic as "repetition improves memory" is "either invalid or needs qualifying" (p.228). Roediger's point is not that replicable patterns of memory behaviour do not exist, but rather that such patterns are always highly specific to a very narrow range of circumstances. What is lacking is any kind of universal or truly generalizable laws of memory (hence the universality of the qualification, "it depends").

Even the plausibly simpler task categorising types of memory has proven quite difficult. In his eloquent paper, Tulving (2007) reveals the vast, confused landscape of memory taxonomy in the contemporary literature. We have already seen, in relation to attention and consciousness, that brains are messy, mushy things, and their activities tend to defy neat classification. Even what qualifies as a memory can be unclear, as in the case of *Rilkean* memory—autobiographical memory that is neither episodic nor semantic—things like the emotional response one has to certain stimuli without consciously being aware of the past experiences that are the reason for that response (Rowlands, 2015). Classifications of different kinds of

¹⁰⁷ *Theaetetus* 191c-e.

memory, then, often bear little relevance to each other, or can pick out different aspects of the same process. Michaelian and Sutton (2017, sec. 2) suggest that there is an emerging consensus on the taxonomy advocated by Squire (2004, 2009), which divides memory into *declarative* memory, which represents the world and aims at truth—e.g. remembering *that* an event happened (including *episodic* or recollective memory and *semantic* or propositional memory), and *non-declarative* memory, which is a definition by exclusion (whatever is remembered but is not Declarative) and which does not represent the world or aim at truth—e.g. remembering *how* to do something. But this distinction seems quite independent of the distinction, say, between *implicit* (subconscious) and *explicit* (conscious) memories, or between *direct* memory, which comes to a subject automatically, and *generative* memory, which takes an effort to recall.

One of the memory concepts that has emerged and proved popular in recent times is *working memory* (WM). It developed out of the older distinction between what we might call *short-term memory* (STM) and *long-term memory* (LTM), although those technical terms are somewhat late arrivals.¹⁰⁸ Intuitively, we all know the difference between remembering, say, a phone number for seconds to minutes, then forgetting it, and remembering a phone number more or less permanently. In the days before writing emerged, many cultures developed quite sophisticated practices for transferring information from STM to LTM, often resulting in quite prodigious feats of memorising huge narratives with stunning accuracy. Today, this art is practiced using ancient techniques such as the “mind palace” and is the subject of intense international-level competition.¹⁰⁹

While the concept of LTM remains quite viable today—in that it seems to capture a genuinely distinct and relatively well-circumscribed cognitive ability—the concept

¹⁰⁸ Some other terms of historical interest (although they still pop up in the contemporary literature) include *primary* (=STM) v *secondary* (LTM) memory, due to James (1890), and *immediate* memory (STM), which first enjoyed a vogue in the 1880s (Ackerman, Beier, & Boyle, 2005, p. 31).

¹⁰⁹ See for example, Pillai (2016). The World Memory Sports Council has conducted an international memory competition since 1991, see <http://www.worldmemorychampionships.com>.

of STM has fallen out of favour in many circles (though not all) in recent decades.¹¹⁰ Unlike LTM, STM has come to be seen as more of a suite of distinct abilities, such as iconic memory, fragile visual short term memory, and WM.¹¹¹ The connection between WM and the earlier concept of STM is therefore in some ways one of supersession: “The term ‘working memory’ evolved from the earlier concept of STM, and the two are still on occasion used interchangeably” (Baddeley, 2012, p. 3).¹¹² It is a distinction that many authors still find of scientific interest,¹¹³ and the term STM has certainly not disappeared over the historical horizon.

5.3 What is Working Memory?

Since it arose, the concept of WM has been remarkably pervasive. As Baddeley (2012, p. 2) notes, his recent review of working memory (Baddeley, 2007) contained 50 pages of references, and according to Logie and Cowan (2015, p. 315) Baddeley and Hitch’s (1974) original chapter has been cited more than 10,000 times in the

¹¹⁰ For an recent account of the demise of STM, see Danziger (2008, pp. 176–182). For an earlier account, see Crowder (1982). For an account of the role of *measurement* of STM in the demise of the *concept* of STM, see Richardson (2007).

¹¹¹ For a nice account of the history of the fractionation of memory generally, and its empirical foundation, see Squire (2004). For the history of the more specific relationship between LTM, STM, and WM, see Cowan (2008).

¹¹² The first use of the term “working memory” was in reference to computer memory in Newell and Simon (1956), which inspired its first use in human cognition by Miller et al., (1960). Logie & Cowan (2015, p. 320) provide an overview of ideas that led up to WM, from Locke, through Wundt, James and Broadbent. For more accounts of the history of WM, see Baddeley (2003, 2010, 2012), Postle (2006), Repovš & Bresjanac (2006), and Soto & Silvanto (2014).

¹¹³ See for example, the discussion of the distinction between STM and WM in relation to intelligence in Ackerman et al., (2005).

scientific literature, justifying their assessment of it as being “one of the most influential works in the field of cognitive science.” Baars and Franklin (2003) among others, agree with this assessment: “Baddeley and Hitch’s working memory model is probably the most influential integrative model of cognition of the last few decades” (p.166). But the cost of such pervasiveness is that—like attention and consciousness—the concept of WM has been reworked, reimagined, and reused in a bewildering range of ways.¹¹⁴

“Once there was a short-term store—a system responsible for the memorization of a small number of chunks for the time one needs to walk from the phone book to the telephone. Over the past three decades, this system has evolved into the central stage of higher-order cognition. Now called working memory, it has been associated with an increasing number of basic cognitive functions, up to a point where it sometimes appears as a conceptual ragbag for everything that is needed for successful reasoning, decision making, and action planning” (Oberauer, Süß, Wilhelm, & Wittman, 2003, p. 167).

In this treatise I will bracket this ongoing debate over how best to capture that basic cognitive function of holding content for brief periods of time and manipulating it in useful, often goal-oriented ways. Like attention, WM seems to melt into many other cognitive processes (Cowan, 1999, p. 63). But models of WM provide enough clarity of concept to allow us to achieve what is necessary for my arguments in Part II—a principled way of demarcating Executive Attention (rather than Liberal Attention) and of thereby identifying its presence or absence, and analysing its relationship to

¹¹⁴ For a discussion of the many uses of the term WM, see Cowan (2017) and Poldrack and Yarkoni (2016, pp. 600–601). This is partially due to the *toothbrush problem*: “Psychologists treat other peoples’ theories like toothbrushes — no self-respecting person wants to use anyone else’s” (Mischel, 2008).

consciousness. To that end, I consider the two chief classes of models of WM, multicomponent models and embedded process models.

5.3.1 Multicomponent Models

Working memory is:

“a hypothetical limited-capacity system that provides the temporary storage and manipulation of information that is necessary for performing a wide range of cognitive activities” (Baddeley, 2012, p. 7).

The author of these words is Alan Baddeley, who, together with Graham Hitch, introduced the Baddeley Hitch model of WM (henceforth, “BH”) that has dominated the field since the early 1970s (Baddeley & Hitch, 1974). Too many authors to mention have taken this basic model and developed it in interesting ways. Baddeley himself identifies four alternative models to BH (Baddeley, 2012, pp. 19–22). However, for my purposes, I will divide currently proposed models of WM into just two categories that I will now describe: *multicomponent models*, such as BH and its intellectual descendants, and *embedded models*, such as those of Cowan (1988, 1999) and Postle (2006).¹¹⁵ I have tabulated a range of variations and developments of these two categories (Appendix 5), as well as some common metaphors (Appendix 6), some of which I refer to in the discussions below.

¹¹⁵ This distinction could be seen as an application of the difference between *systems views* and *process views* of memory, see for example Bechtel (2001).

Initially, the BH multicomponent model included only three components. The *visuo-spatial sketchpad* stores visual and spatial information for short periods of time and makes them available for manipulation by the central executive, and the *phonological loop* similarly stores auditory information. The *central executive* subserves the manipulation function. This is described by Baddeley (1996) as a kind of homunculus, a place holder in need of further elaboration. But its roles are fairly well defined, and include, though they are not limited to, access to LTM, attentional control of action, and conscious access to relevant contents of the other components. Later, a fourth component, the *episodic buffer*, was added to account for integrated multimodal representations, which could not be accommodated by the modality-specific visuo-spatial sketchpad and phonological loop (Figure 4A, below). The central executive, interacting with the other components, is thus eminently suited to demarcate the locus of Executive Attention as distinct from Liberal Attention.

The multicomponential nature of BH grew out of empirical data: a dissociation between the effects of phonological and visual distractors. Things that impeded visual WM seemed to have little effect on phonological WM operating at the same time, and vice versa. However, BH was always meant to be a concept rather than a detailed working model: “My overall view of WM therefore comprised, and still comprises, a relatively loose theoretical framework rather than a precise model that allows specific predictions” (Baddeley, 2012, p. 7). In the same paper, he speculates about ways other modalities—and indeed the rest of cognition—might be integrated into the core of BH (see especially his Figures 4 and 5).¹¹⁶ This raises the first pertinent question:

¹¹⁶ The *Interactive Cognitive Subsystems* Model of Barnard (1999) attempts to include all modalities. Baddeley (2012, p. 21) observes that it “can also be mapped directly onto” BH. But it also departs from BH in some important ways. For example, there is a noticeable absence of any central executive in the model, but “central executive functions are themselves accomplished by processing interactions among subsystems” (p. 298).

5.3.1.1 The Fractionation Question

Is there just one WM system encompassing all modalities / cognitive functions, or are there several systems?

Relevance: e.g., the relationship between attention and consciousness might be teased apart by considering how each relates to different components of WM.¹¹⁷

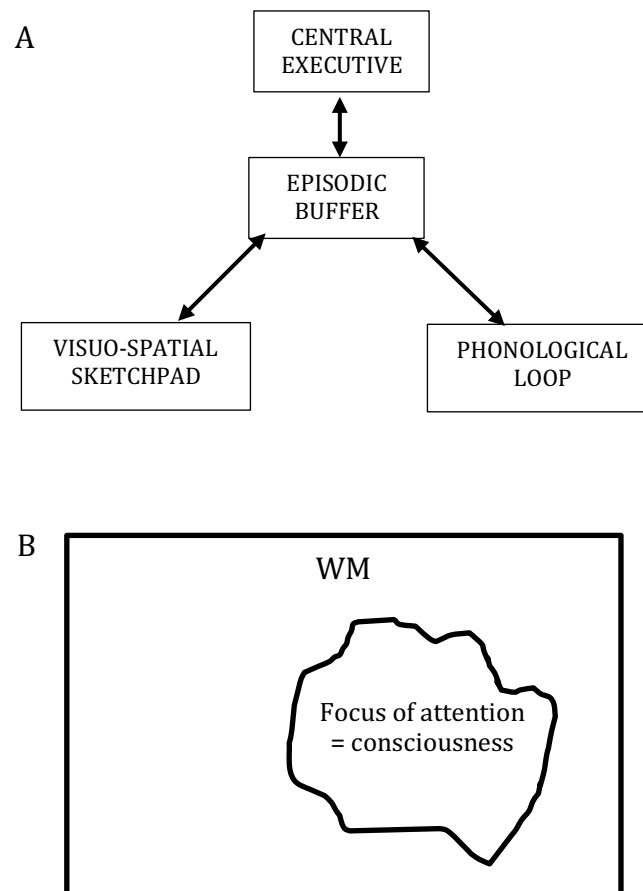


Figure 4. A. BH's multicomponent WM. See Figures 1 and 5 in Baddeley (2012) for how this model has developed. B. Cowan's (1988, 1999) embedded processes model of WM.

¹¹⁷ I return to the Fractionation Question in Chapter 8 (2.1).

A related issue that has been of significant empirical interest is whether there is such a thing as *implicit* WM, WM activity that is subconscious.

5.3.1.2 *The Subliminal Question*

Is there a subconscious or implicit WM?

Relevance: if there is indeed such a thing as implicit WM can it be attended? This promises to be an arena in which we can untangle the relationship between attention and consciousness.

Hassin et al., (Hassin, 2013; Hassin et al., 2009), Hsieh and Colas (2012), Samaha et al., (2016) and many others have argued powerfully for implicit WM, while Prinz (2012) has argued against such conclusions. If there is indeed such a thing as implicit WM, the question arises as to whether attentional strategies are involved, which would constitute a case of *A~C* (see 6.2).

5.3.2 Embedded Models

Embedded models (e.g., Cowan, 1988; Postle, 2006) see WM as something that happens in all or most cognitive processes, rather than being a distinct system with its own structure and components. For Postle,

“working memory functions are produced when attention is directed to systems that have evolved to accomplish sensory-,

representation-, or action-related functions. From this perspective, working memory may simply be a property that emerges from a nervous system that is capable of representing many different kinds of information, and that is endowed with flexibly deployable attention. Predictions about the nature of representations contributing to the short-term retention of any particular kind of information are made by considering the nature of the information that is to be remembered, and the mental processes that are afforded by the task that is being performed” (Postle, 2006, p. 29).

Embedded models (Figure 4B, above) strongly imply a widely distributed neural implementation, as this kind of implementation reflects the function.¹¹⁸ Embedded models see WM as being inextricably enmeshed in other cognitive processes, or perhaps even arising or emerging (in Postle’s words) from the overall cognitive economy, or parts thereof. Thus, on embedded models, there is no discrete WM store—only something like Cowan’s idea of “activated LTM,” where the content being processed by WM is *just* LTM traces that have been activated in a particular way.¹¹⁹ This raises another two pertinent questions:

5.3.2.1 *The Duplication Question*

Are contents of perception or memories retrieved from LTM somehow duplicated in a distinct WM store?

Relevance: it is important to understand whether one attends to or becomes conscious of information only within WM or also outside

¹¹⁸ For a distributed account of WM, see Christophel et al., (2017). For a recent review of the neural correlates of WM and their relation to Q, see Fazekas and Nemeth (2018, pp. 3, 5–6).

¹¹⁹ For a more nuanced embedded model, see Oberauer (2002). For an argument against activated LTM models, see Norris (2017).

of WM. This may be one area where attention and consciousness behave differently.¹²⁰

5.3.2.2 *The Manipulation Question*

What exactly does it mean when say that information is being manipulated? How does storage relate to manipulation?

Relevance: how do the functional roles of attention and consciousness in manipulation illuminate their relationship to each other? E.g., do they play distinct functional roles?

One of the models that addresses this question is Jonides et al's (2008) *Mind and Brain Model*. The complexity of their model makes it difficult to summarise briefly (see pp. 212-213 for their own summary), but of particular interest here is the important role played by *attention*. In brief, they characterise the information to be manipulated as being in a “dormant state,” whether in LTM or in perceptual stores, until activated. The “focus of attention” is then just those activated traces that are accessible for “cognitive action.” They make some further points about how such states might be maintained in, or lost from, the focus of attention. This model makes sense of the intricate connections between representations—the “context” of each memory or perceptual trace—which clearly play a part in how that trace behaves. I consider the Manipulation Question in a little more detail below (5.4.2).

Standing outside the multicomponent / embedded distinction, but cutting across it, is the *computational modelling* approach, which can be applied to both kinds of model. Miyake and Shah (1999) present a number of very detailed attempts at

¹²⁰ I address this question 5.4.1, 7.2.1.3, and 9.3.2). One account of actual neural duplication of traces is that of Jacobs and Silvanto (2015).

computational modelling of WM. While such approaches have their benefits (e.g., Lovett, Reder, & Lebriere, 1999, p. 136), they are extremely complex, and susceptible to the objections that have been raised against computationalism generally (e.g., Rescorla, 2017, sec. 7).

Inherent in nearly all models of WM is the idea—based on a large body of empirical work—of only a limited amount of content being capable of being maintained and manipulated in complex ways at any one time:

5.3.2.3 *The Capacity Question*

What is the magnitude and the nature of the capacity limitations of WM?

Relevance: the capacity limitations of attention and consciousness are interesting questions in their own right currently under investigation. Yet more empirical work has been conducted into the capacity limitations of WM. This empirical work may fruitfully be brought to bear on the other two concepts, thus helping to illuminate their relationship.¹²¹

Both approaches deal with the question of capacity limitations. On BH, the central executive-episodic buffer axis is of limited capacity, while on Cowan, it is the focus of attention that limits capacity. The emphasis in embedded models is on understanding WM, not by teasing out the functions of its components, but by teasing out how it functions within the broader scheme of cognition, which raises another pertinent question:

¹²¹ I address this question in substantial detail in Chapter 8.

5.3.2.4 *The Integration Question*

Where does WM fit, and how does WM interact with the rest of cognition?

Relevance: attention and consciousness are both global cognitive processes. WM promises to help elucidate the roles they play in the cognitive economy with a finer resolution, and therefore possibly help shed light on the relationship between them.

And of course, a specific sub-question of the Integration Question is how WM can help us to make progress on Q:

5.3.2.5 *The Correlation Question*

What is the relationship between WM, attention, and consciousness, and in what ways do they correlate with each other?

Relevance: WM's intimate connection to both attention and consciousness makes it a valuable tool for investigating Q.

The difference between the BH and Cowan is perhaps more a difference in emphasis than anything more substantive. Multicomponent models tend to emphasise a broadly modular structure to WM, although Baddeley (2012, p. 7) goes to some pains to explain that his components are certainly not Fodorian modules (Fodor, 1983), in that they are much less encapsulated and more integrated with each other and with other non-WM systems. Burgess and Hitch (2005) propose a model of the BH phonological loop that connects it intimately to LTM, raising the question of just where the line ought to be drawn between WM and LTM. One might view their

model as multicomponent, since it is explicitly founded on the BH components, yet it also has features of embedded models in that the functions of WM seem to be tightly entwined with LTM. Nonetheless, the componential character of BH is central to it—it is by understanding the components and the functions of the components that we will build an understanding of WM. The question of whether the components are anatomically distinct entities is of course left open. They may be implemented neurally in either a localised or a widely distributed fashion. This is a question to be answered empirically, and has little bearing on the usefulness of the model as a description of how WM *functions*. It also raises the last of our pertinent questions:

5.3.2.6 *The Neural Question*

*How is WM implemented in the brain?*¹²²

Relevance: a great deal of work has been done on the neural correlates of consciousness, and to a lesser extent, on those of attention. Identifying the neural correlates of WM should help shed light on both, and therefore on the relationship between attention and consciousness.

¹²² For an exploration of the plethora of possible neural mechanisms underlying WM see Barak and Tsodyks (2014) and Serences (2016). Michaelian & Sutton (2017, sec. 5.2.1) address the question of whether philosophers ought to engage in neural questions like this. The short answer is, yes. For a recent presentation of two opposing views, see Lundqvist et al., (2018) and Constantinidis et al., (2018).

5.3.3 Metaphors

Models like BH, Cowan, and others are not the only way to understand WM. Plato's metaphor of the wax tablet (5.2) is an early example of the time-honoured practice of understanding by analogy or metaphor (three of which I describe in this section), which has a rich history in enquiry into memory.

“Throughout its history, memory discourse has provided a rich field for the play of metaphors. This continued to be the case even after memory became a topic for scientific psychology. In fact, this area of psychology is unusual in the frankness with which the role of metaphor has been widely recognized. Little more than a decade ago a discussion of metaphors in memory research in the journal *Behavioral and Brain Sciences* drew in some twenty-five contributors, the great majority of them experimental psychologists. Yet well over two thousand years ago metaphor already played a major role in the first sustained discussion of memory in Europe, that of Plato. Nor is it difficult to find numerous examples of memory metaphors during the intervening centuries” (Danziger, 2008, p. 24).

Unlike a model, a metaphor does not seek to simulate as closely as possible the reality being explained, but rather to draw a connection between central or salient features common to both reality and metaphor. The hope is that, by recognising this similarity, the complex target is better understood, since it is “like” the more familiar metaphor. It is important to note that the mechanisms involved in producing the behaviour of the metaphor and the target need not be identical, or even similar—only the results need to be similar in relevant ways. It is thus dangerous to take metaphors too far, to extrapolate unwarranted conclusions from them about the

reality.¹²³ Metaphors help us understand the nature of WM and play a role in my arguments in Part II, so below, I describe three classes of WM metaphors: *map rooms*; *boxes*; and *manipulation metaphors*, tabulated in Appendix 6.

5.3.3.1 *Map Room Metaphors*

Well before the term WM was coined, people were thinking about how we go about using the information we have to think and act.

“We assert that the central office itself is far more like a map control room than it is like an old-fashioned telephone exchange. The stimuli, which are allowed in, are not connected by just simple one-to-one switches to the outgoing responses. Rather, the incoming impulses are usually worked over and elaborated in the central control room into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release” (Tolman, 1948, p. 192).

The metaphor of the *telephone exchange* is clearly inadequate. Our minds integrate information in ways far more intricate and complex than simple one-to-one connections between discrete inputs and outputs. But what exactly does the *map room metaphor* mean? What exactly is the map? On what canvas is it drawn (Duplication Question), and where is this canvas (Neural Question)? Need there be conscious control directing the outputs (Subliminal Question)? How does

¹²³ For further discussion of metaphors of memory generally, see Part I of Groes (2016).

information on the map come to be related in law-like ways to both inputs and outputs (Manipulation Question, Integration Question)?

5.3.3.2 *Box Metaphors*

Plato's wax tablet metaphor is one member of a class of metaphors I will call *box metaphors*. Here, WM is a box into which memories are placed and removed as needed. The limited size of the box imposes the required limitations on the capacity of WM (Capacity Question), and the idea of some subject focusing attention and being aware only of what is in the box captures nicely the close relationship between WM and attention and consciousness (Correlation Question). Different kinds of objects can be placed in the box in different relations to one another (Fractionation Question).

Variations on the box metaphor include memory (generally, not just WM) as a storehouse, a repository, an archive, or more poetically as "the stomach of the soul" (Augustine) or the "repository of forces" (Hobbes).¹²⁴ A more interesting variation on this class of metaphors is to see WM as a working desktop.¹²⁵ One removes books from a shelf and places them on a physical desktop, but the files on a computer's hard drive are duplicated in working memory and represented on the computer's desktop while the original files remain intact *in situ* (Duplication Question). Both kinds of desktop exhibit limited capacity. Postle (2006) characterises the "standard model" of WM (likely referring to BH and its descendants) as being something like this desktop metaphor. Embedded process models are more like a reader picking up a book off the shelf and reading it there, rather than taking it back to a desktop.

¹²⁴ Malyshkin (2013, p. 38).

¹²⁵ My own invention.

One of the things the map room and box metaphors imply is that there is *someone* in the map room or library, looking at and manipulating the information on the map or the items in the box, taking note of new information or items and making decisions and acting accordingly. On Baddeley's account, the central executive can play this role of the metaphorical homunculus¹²⁶—the controller of the map room or the reader in the library. In fact, it is possible to read Baddeley's WM as playing the role of Baars' global workspace:

“The central core of working memory would incorporate and synthesize information from many different modalities and codes (both sensory and semantic) to arrive at an overall conception of the environment and of one's current situation” (Logie & Cowan, 2015, p. 318).

I therefore use BH's central executive and Baars' global workspace as roughly interchangeable concepts since they play the same role relevant to my discussions of Q: the personal-level coordination of cognition,¹²⁷ which plays an integral role in what we normally think of as our ongoing cognitive life.

5.3.3.3 *Manipulation Metaphors*

A third class of WM metaphors focuses on the way content is manipulated in WM.

¹²⁶ For a discussion of the homunculus problem, see Ramsey (2007).

¹²⁷ See also Baars and Franklin (2003) for a discussion of how BH might be harmonised with Baars' Global Workspace. Not all accounts hold the BH central executive to be equivalent to Baars' global workspace (8.6.2.3).

“Adding to the confusion is that a number of different metaphors are used to refer to working memory, and to highlight different characteristics of the concept, including the “box” or “place” metaphor, the “workspace” or “blackboard” metaphor, the “mental energy” or “resources” metaphor, and the “juggling” metaphor” (Miyake & Shah, 1999, p. 2).

The last three metaphors mentioned above constitute a third class of metaphors in addition to box metaphors and map room metaphors—*manipulation metaphors*. The *juggling* metaphor captures the limited capacity to skilfully manipulate only a certain number of objects, as well as the constant danger of losing “concentration” and seeing them all tumble to the ground, an experience anyone who dabbles in challenging IQ puzzles knows only too well. Similarly, the idea of a limited quantity of mental energy or resources being available for manipulation captures the same idea. This is a very neat way to capture and perhaps begin to explain the empirically observed limitations of WM. It relates nicely to models of *attention* that are based on competition, such as that of Ruff (2011), which accommodates that close connection between WM and attention we have so often observed. Box metaphors emphasise the storage of content, while manipulation metaphors emphasise how content is processed. This duality of function in WM is so central, it is even reflected in the term: “working” and “memory,” and is worth exploring a little further.

5.4 The Dual-Aspect Nature of Working Memory

A strikingly consistent feature of all models and metaphors is that WM serves *two* chief functions: the *short-term storage* of information; and the *manipulation* of this

information (e.g., Cowan et al., 2005).¹²⁸ This manipulation aspect of WM sets it apart from other forms of memory (e.g. iconic memory, LTM) which are usually conceived of as serving only the storage function, and merely providing the content for other cognitive systems to manipulate. It also led Baddeley to ponder whether a better term might have been “working attention” (Baddeley, 1993b, pp. 167–168),¹²⁹ but as I observed earlier, “working memory” does capture the dual functions nicely. In this section, I elaborate on these two themes a little more and use them specifically to delineate a sense of Executive Attention employed in my investigation of Q in Part II.

5.4.1 Storage

WM is said to include a short-term, limited capacity store that maintains content for moderate periods of time, making it easily accessible to manipulation. The Duplication Question raised one way this might occur: content may be “copied” from another location—say LTM, or temporary content in the visual cortex—to a discrete and dedicated WM storehouse elsewhere in the brain. Call this the *Duplicate Store Model*. But this seems profligate. The brain has evolved to function efficiently: “we know that neural tissue and activity are expensive; see Aiello and Wheeler (1995); and we also know that as a result of such constraints, the wiring diagram for the brain is about as efficient as it is possible for it to be; see Cherniak et al. (2004)” (Carruthers, 2016, sec. 4). Duplication of identical content in multiple sites seems to

¹²⁸ This functional duality in memory is not new. Seneca recognised that memories are not simply stored, but interact with each other and evolve: “We should imitate bees, we should mingle all the various nectars we have tasted, and then turn them into a single sweet substance, in such a way that, even if it is apparent where it originated, it appears quite different from what it was in its original state,” quoted in Carr (2010, p. 122). Fazekas (2018, p. 6) adds a third function of *monitoring* of content in WM, but this plausibly falls under the general function of manipulation.

¹²⁹ See also Beaman (2010).

be a waste of resources, both in terms of neural “real estate” and energy consumption. We are still trying to understand the neuronal basis of content transfer and storage,¹³⁰ but the problem is made worse by the requirement for then amending the LTM trace to conform to the more processed trace in WM, a process likely repeated over and over. What is more, this model suffers from all the problems identified above with box metaphors of WM storage (5.3.3.2).

Some of these problems with the Duplicate Store Model might be alleviated if we posit that it only duplicates information from elsewhere in the brain partially, only the bits really needed for the project at hand. Call this the *Partial Duplicate Store Model*. But this raises another problem: how is the right set of content chosen for duplication in WM? This is the role Prinz (2012) suggests for attention—the gateway to WM. But while partial duplication alleviates some of the profligacy concerns, it does not eliminate them.

A much more parsimonious model is what I’ll call the *File Directory Model*, taking a cue in the true spirit of WM research from computer science. This model bears an affinity to embedded models, with WM storing only dynamically evolving list of *pointers* to content in other stores, rather than duplicating the content itself. This list defines what is currently being manipulated by, or is easily accessible to, the executive or focus of attention of WM. This model is not only more parsimonious, but also in harmony with how we understand memory retrieval to work neurally.¹³¹ It could plausibly be the mechanism of WM access to perceptual stores. And it solves the problem of needing to re-encode refined content into LTM—the LTM traces themselves, *in situ*, are the things being refined.

A fourth possibility is to eschew storage in a discrete WM store of any kind, even pointers. On some embedded models, WM is *just* activated LTM, or perceptually, an increase in connectivity between, say, iconic memory and WM (Landman,

¹³⁰ E.g., how spatial LTM is encoded and accessed (Jadhav, Kemere, German, & Frank, 2012).

¹³¹ This is still an area of intense investigation. See Jonides et al., (2008) and the articles in *Neuroscience* issue 139 (2006).

Spekreijse, & Lamme, 2003, p. 163). In other words, there is no discrete WM store at all, not even a list of pointers. Call this the *No Store Model*. It is the most parsimonious and efficient of the four models, but at the cost of calling into question the very concept of WM in the first place. If there is no unique *store*, why call it “memory”? Why not just employ the concept of an executive, global workspace, or Executive Attention alone? But this would be to ignore the copious empirical evidence for the maintenance in some way of a subset of content that is capable of being manipulated in complex ways. If the exact model of WM storage is still unclear to us, the reality of such storage remains compelling.

5.4.2 Manipulation: Executive Attention in Working Memory

None of our four options for understanding the storage aspect of WM seems without complications. I return to these issues in Part II. What of the manipulation aspect? What are we actually talking about when we speak of stored information being manipulated (Manipulation Question)? Here too there seem to be two interesting options. The first is an *adverbial* interpretation of WM. Mole (2011a) argues that some cognitive processes can be characterised by the predicate “attended” while others cannot. Attention is not a noun, is not a process, but a characteristic of some cognitive processes. Similar arguments have been made for consciousness. Might we interpret WM in the same way? Thus, some cognitive processes would be characterised by being “WM’ed,” while others would not.

One possible problem here might be how to explain the high degree of integration and coordination that is undoubtedly characteristic of WM. Perhaps such coordination emerges naturally from the operations of many divergent cognitive functions, but it would certainly be more likely that there is indeed some kind of central executive as in almost all models and metaphors of WM (Appendices 5 and 6), controlling the flow of processing of content in order to achieve higher-order

outcomes. This picture provides a principled way for delineating the concept of Executive Attention I distinguished from Liberal Attention in 3.4.4.

Simply, *Executive Attention can be taken to be all and only the implementation of attentional strategies by the cognitive executive*. This is the characterisation of Executive Attention I employ to explore Q in Part II. It is a characterisation that has found proponents elsewhere. For example, “Ned Block, however, has in most recent publications interpreted his own concept of access as being identical to the contents of ‘working memory’ (Block, 2011a; Carruthers, 2017)” (Mogensen & Overgaard, 2018, p. 1). Thus also Baddeley (1993b, pp. 167–168) points out that the central executive does not itself perform any storage tasks, only manipulation, and considers the central executive to be “primarily attentional in nature.” Others have also identified “Executive Attention”¹³² with the attentional strategies implemented by the cognitive executive (Engle et al., 1999, pp. 104–105; Kane & Engle, 2002, p. 638).

5.5 Boundaries of Working Memory?

The involvement of WM in quite diverse cognitive functions (multimodal perception, motor control, etc.) threatens to bloat our concept of WM until it encompasses all cognition, much as Liberal Attention does. Again, that would make WM a far less useful concept. Fortunately, both the multicomponent and embedded processes models avoid this trap. This is obvious from the characteristics of the central executive-episodic buffer axis (BH) or the focus of attention (Cowan)—these are empirically discernible and measurable systems that are specialised, of limited

¹³² Kane and Engle’s definition of Executive Attention differs from mine in that it involves only “controlled” (top-down) attention.

capacity, and manipulate selected content in specific, coordinated, goal-oriented ways, while an awful lot of other cognitive processing is going on unattended and unconscious “in the background,” so to speak.

Another way to see this is to consider the storage and manipulation functions of WM. In terms of *storage*, there is clearly a difference between WM stores and, say, LTM or iconic memory stores. WM stores are demonstrably far more limited in capacity than the stores of either LTM or iconic memory, and while iconic memory stores are retained for much briefer periods of time than WM stores, LTM stores are retained for much longer periods. Clearly, an awful lot of content is stored outside of WM. Similarly, when we turn to *manipulation*, we see that an awful lot of content is manipulated outside WM, only some of which ultimately comes to be manipulated by WM, precisely that which is, on my definition, Executively Attended. The early visual processing that binds content together into objects, for example, is complex manipulation that proceeds independent of WM.¹³³

5.6 Chapter Summary

WM is an empirically-derived construct that frequently intrudes into discussions of the relationship between attention and consciousness. In this chapter, I have surveyed briefly its history and the models and metaphors that have been employed to capture its nature, and found that the dual nature of WM—storage and manipulation—connect it closely to attention, consciousness, and the rest of the cognitive economy. This survey raised a number of fascinating pertinent questions—none of which are original—applicable to WM and to other cognitive

¹³³ I argue for the distinctness of Executive Attention and binding in 3.4.2. For an overview of the neuroarchitectural integration of WM with the rest of cognition, see Eriksson et al., (2015).

concepts, which serve to scaffold investigation into the nature of cognition generally, and our quest to answer Q specifically. The manipulation aspect of the central executive (BH) or focus of attention (Cowan) of WM in particular gives the concept of Executive Attention a firm foundation and clearly demarcates it from attentional strategies employed elsewhere in the cognitive economy (Liberal Attention), providing a principled basis particularly for my discussion of overflowing phenomenal content (Chapter 7) and capacity limitations (Chapter 8).

5.6.1 Conclusion to Part I

Part I has exhaustively staked out the conceivable territory of patterns of relationships between attention and consciousness (STF) and the four kinds of relationship (Chapter 4), and narrowed down an interesting and plausible set of definitions for the concepts of consciousness (Chapter 2), attention (Chapter 3), and WM (Chapter 5). On this foundation we have now at the tools with which to attempt an answer to a specific version of Q in Part II.

PART II

Addressing Q.

6 Pulling Attention and Consciousness Apart

6.1 Recap and Outlines

The goal of this thesis is to address Q: *Is attention both necessary and sufficient for consciousness?* In the broad literature that touches upon Q, authors approach different issues in different ways using different definitions to answer different questions. In Part I, I laid some foundations to try to formulate a specific question (Q) to be answered in a specific way, using specific definitions of attention and consciousness, and employing a specific idea of WM. My task in Part II is to attempt an answer to Q, so posed: *“Is Executive Attention both necessary and sufficient for phenomenal consciousness.”* The terms are defined thus:

Consciousness: *phenomenal experience, characterised by four core features: “what it is like-ness;” a situated first-person perspective; unity; and temporality.*

Attention: *a suite of multiply realisable strategies for structuring—in the human case—cognitions by selection for further processing. Executive Attention is the special case of attention, implemented by the cognitive executive, that much of the literature on Q means by the term “attention.”*

Working Memory (WM): *short-term, limited capacity storage and manipulation of information, described by two important models of WM: the BH multicomponent model, and Cowan’s embedded process model.*

This version of Q is particularly interesting and consequential because all three of these concepts provides an intuitive answer to the question, “what is my ‘self?’” My stream of phenomenal experience and my chain of implementation of Executive Attentional strategies in WM are both highly plausibly ways of capturing that which I take to be my mental “self.” Pre-theoretically, they ought therefore to coincide perfectly. If Executive Attention and phenomenal consciousness *can* come apart, I am forced to re-evaluate the very nature of my “self.”

6.1.1 Outline of Part II

Conscious attention is what we take most of our daily experience to be, while the cognitive economy is numerically speaking largely a black economy, with most of the processing occurring in the absence of either Executive Attention or consciousness. Since the Combinations $A\&C$ and $\sim A\sim C$ are widely accepted as being

instantiated in human cognitive economies,¹³⁴ which of the four live STF Scenarios obtains must be decided by the status of the remaining Combinations: $A \sim C$ and $C \sim A$. My path, therefore, to finding an answer to Q in Part II looks like this: in this chapter, I explore putative empirical cases of $A \sim C$ and $C \sim A$. While there is convincing evidence for $A \sim C$, $C \sim A$ is more controversial, and I devote the bulk of my enquiry to it, but fail to find a compelling candidate. In Chapter 7 I turn to the most promising candidate for $C \sim A$, phenomenal overflow, and build a strong abductive case for it in one particular condition—foveal vision. In Chapter 8 I strengthen the case by building an abductive argument from the relative capacity limitations of attention, consciousness, and WM. I conclude this treatise in Chapter 10 by summarising various answers to various forms of Q, reflecting upon territory covered, and drawing out some interesting implications, consequences, and applications.

6.1.2 Outline of this Chapter

In this chapter I search for cases of $A \sim C$ and $C \sim A$. I begin each search with some reflections on what each Combination means in a human cognitive economy, and then turn to a systematic search, not only of the experimental literature, but also ranging across normal human experience. Some of the candidates I consider are (as far as I know) novel in the context of Q. My search for $C \sim A$ ranges over fifteen different candidates (and others considered in less detail) but is structured in a novel framework of five different ways there might be $C \sim A$, given the definitions of the terms above. None of these are conclusive, but a sixteenth—phenomenal overflow (Chapter 7)—is more promising.

¹³⁴ For numerous arguments and empirical examples of $A \& C$ see Prinz (2011), and Koch and Tsuchiya (2007, p. 17), and for $\sim A \sim C$ see Mack and Rock (1998), and Koch and Tsuchiya (2007, p. 17).

6.2 Attention Without Consciousness

What would it mean to find $A \sim C$, and are there any such cases in the human cognitive economy? These are the questions I address in this section. I begin by addressing some preliminary analytical and methodological issues before turning to empirical evidence for $A \sim C$.

6.2.1 Preliminaries

The *presence of attention* is relatively easy to identify by report, behaviourally, or even by neural signatures. When a subject reports that she has attended a target, the very act of reporting itself is mediated by Executive Attention, and therefore confirms the presence of attention. The same may be said of behaviours such as pressing a given button upon noticing a target. And the N2pc component is an ERP signature of attention (Eimer, 1996; Luck & Hillyard, 1994). All of these are reliable indicators of the presence of Executive Attention.

It is more difficult to identify an *absence of consciousness*. Unlike attention, report is no guarantee that a subject was not conscious of a target. The same may be said of behaviours, and the neural correlates of consciousness are still unknown. I discuss these difficulties at some length in 7.4¹³⁵ and come to some unexpected conclusions in 9.4.3, but for now, I take the empirical studies below at their word, and assume

¹³⁵ See also Simons et al., (2007).

that they have a reasonably reliable and coherent way of identifying that consciousness is absent, unless otherwise stated.

Finally, some of the candidates for both $A \sim C$ and $C \sim A$ will be *local* Combinations. For example, attending to a target in the periphery of vision without being aware of it is local $A \sim C$, since in other parts of the visual field there is $A \& C$. Other candidates will be *global* Combinations. For example, complete sensory deprivation (6.3.3.2) is putatively global $C \sim A$, since there is a putative temporary absence of Executive Attention throughout the whole cognitive economy. It is important in each case to keep attention and consciousness comparable—not to correlate, for example, global attention with a local lack of consciousness.¹³⁶ A global instance of a Combination like $C \sim A$ is not the same thing as a Scenario. A Scenario describes the ongoing cognitive economy that is the very nature of the organism *over time*, whereas the putative $C \sim A$ global Combination of sensory deprivation is a temporary and relatively rare state for a subject.

6.2.2 Empirical Candidates for $A \sim C$

All that is needed to establish $A \sim C$ as a Combination instantiated in the human cognitive economy is just one clear case of any kind of Executive Attention without any phenomenal consciousness at all, global or local, and congruent with respect to my four questions from 1.4.6: Target; Timing; Variety; and Consequences (recognisable as such by being capitalised).

¹³⁶ Compare the concepts of *global* or *background* consciousness such as having a waking experience, on the one hand, and *local* or *specific* consciousness such as seeing a face, on the other (Windt et al., 2016, p. 872).

It turns out that the list of potential candidates is a very long one, and I shall only be able to consider a small sample. However, it only takes one conclusive case to establish that $A \sim C$ is instantiated in a cognitive economy. I therefore begin with some illustrative problems with some candidates before moving on to what I consider to be quite conclusive cases. I conclude the section by briefly addressing some general challenges to these kinds of cases, none of which challenges are, to my mind, compelling.

Tononi and Laureys (2009, pp. 378–379) discuss neurological evidence for $A \sim C$. A neural marker of attention, the N2pc event-related potential component, can be elicited normally by stimuli that are not consciously experienced, due to object-substitution masking, which prevents the formation of a phenomenal perception of an object (Woodman & Luck, 2003). Further, Wyart and Tallon-Baudry (2008)¹³⁷ identified two distinct and dissociable MEG patterns corresponding to spatial attention and conscious visual experience respectively. Of course, to confidently assert that one has shown a dissociation between the neural correlates of attention and consciousness, one should first be confident that one has identified the neural correlates of consciousness. Given that is still a long way off (Chalmers, 2000; de Graaf, Hsieh, & Sack, 2012; Koch, Massimini, Boly, & Tononi, 2016),¹³⁸ it might be wise to take this approach with a degree of caution.

Koch And Tsuchiya (2007) consider a number of empirical results that argue for $A \sim C$ (which I called KT2 in 1.3.1). These include: visual crowding; priming; continuous flash suppression; blindsight; and feature-based attention (pp. 17-18). Some of these raise doubts we don't have time to go into here, but among the more convincing cases is that of continuous flash suppression. KT discuss Jiang et al., (2006), in which salient stimuli (erotic images) masked by continuous flash suppression¹³⁹ draw spatial attention to a location in the visual field, thus facilitating the task of

¹³⁷ See also Wyart et al., (2012) and Chica et al., (2012).

¹³⁸ See also an issue devoted to this matter, of *Consciousness and Cognition*, 2017(54).

¹³⁹ I believe that KT mistakenly describe this as continuous flash suppression rather than interocular suppression.

identifying the orientation of a Gabor patch briefly presented subsequently at the same spatial location. From this, KT draw the conclusion that the salient nudes attract spatial attention to their spatial location in the absence of consciousness since the nudes are masked from conscious experience. This is certainly Executive Attention, since the attentional strategies of alerting and orienting (attention, not the eyes) are implemented within the executive process of following the instructions of the experimenters. And the Timing is congruent here, since the nudes are both unconscious and draw spatial attention at the same time. But there is a question about the congruence of the Target: while the subject is unconscious of the *object* (the nude) he is not necessarily unconscious of the *spatial* location in the visual field. And while there is spatial attention to that location, there is no object or feature attention. What would constitute a genuine case of $A \sim C$ here is $A(\text{object}) \sim C(\text{object})$ or $A(\text{spatial}) \sim C(\text{spatial})$. But all we can be sure of here is $A(\text{spatial}) \sim C(\text{object})$. The candidates discussed under KT2 are certainly suggestive, but all suffer from similar concerns.

More conclusive is the work of Kentridge et al., which has been foundational in establishing both $A \sim C$ and $C \sim A$ (Barrett, 2014, p. 13). In favour of $A \sim C$, they showed that cues in the blind region of a blindsight patient's visual field were able to direct the patient's attention to a second region without entering the patient's conscious awareness (Kentridge et al., 1999; Kentridge, Heywood, & Weiskrantz, 2004). If the Target here is taken to be the cue, then the fact that it has the Consequence of directing the subject's attention confirms that the cue itself was Executively Attended, and the fact that the cue is not phenomenally experienced confirms the absence of consciousness of the cue, and both of these are at the same Time. They later generalised the paradigm to the normal population, rather than in blindsight (Kentridge, Nijboer, & Heywood, 2008). In this paper they cite recent work that uses other paradigms (flash suppression, masked priming) that supports their conclusions (Kanai et al., 2006; Sumner, Tsai, Yu, & Nachev, 2006). More recently, they have shown that not only can endogenous, voluntary attention occur without conscious awareness, but exogenous, reflexive attention can also occur without consciousness (L. J. Norman, Heywood, & Kentridge, 2015).

Also conclusive are Soto and Silvanto (2016), who discuss numerous empirical cases of Executive Attention in the absence of consciousness.¹⁴⁰ For example, van Gaal et al., (2012) argue that there is considerable evidence that cognitive control—functions such as “error detection and correction mechanisms, conflict resolution, response inhibition, and task-switching” (p. 1), all of which would normally implement Executive Attention—can be activated and operate upon unconscious stimuli. Here, the Target and Timing conditions are met for *A~C*, since the stimulus being processed by Executive Attention is unconscious, even if the Consequence, the outcome of the process, is conscious.¹⁴¹

In addition to candidates thoroughly discussed in the literature, I venture to suggest some more that have received little attention or none at all in the context of *Q*. I do not present them as conclusive, only as worthy of further exploration.

Proprioception is position sense—the ability to know the state of a skeletal joint, whether your right elbow is flexed or extended, for example, and by how much. Does proprioceptive content enter into WM for processes implementing Executive Attention? That seems indubitable, as in the case of reaching out to pick a cup of tea

¹⁴⁰ See also Soto et al., (2011).

¹⁴¹ Some further strong candidates for *A~C*: middle-level visual content (Jacob, Jacobs, & Silvanto, 2015); cognitive control (Christensen et al., 2016; Lamme, 2018, p. 8); implicit grouping of objects (Kimchi, Devyatko, & Sabary, 2018); object based guidance (Chou & Yeh, 2012); implicit learning during sleep (Andrillon & Kouider, 2016 and volume 122 of *Neurobiology of Learning and Memory*, 2015); limbic status epilepticus (Monaco et al., 2005, p. 156); and priming (Doyen, Klein, Simons, & Cleeremans, 2014). Some controversial candidates: vegetative states with alert-like fMRI activity (Block et al., 2014, p. 556; Klein, 2017; Naccache, 2006, p. 3196); gaze cueing (Kentridge, 2011, pp. 251–252); reading and math (Sklar et al., 2012); experiential blink and its dissociation from attentional blink (Pincham, Bowman, & Szucs, 2016); covert pain and the hidden observer phenomenon (Bitter, 2010; Hilgard, Hilgard, Macdonald, Morgan, & Johnson, 1978); subconscious motivations, memory retrieval, etc., (Kihlstrom, 1997); implicit cognition in hypnosis (Kihlstrom, 2007, pp. 451–460); and implicit social cognition (Greenwald & Banaji, 1995). For general overviews of the research on *A~C*: Kouider and Dehaene (2007); Simons et al., (2007); Kentridge (2011); and Barrett (2014, pp. 11–16).

with your hand, where proprioceptive feedback is essential for the smooth movement that prevents spillage.¹⁴²

Do we phenomenally experience proprioception? Introspection reveals that there is something “it is like” to experience heat on your fingertips,¹⁴³ but there seems to be nothing “it is like” to know *purely* proprioceptively (perceptually rather than propositionally) that your elbow is flexed rather than extended.¹⁴⁴ What we do phenomenally experience about our elbows is not proprioceptive content but visual, tactile, and baresthetic (pressure sense) perceptual content, which must not be confused here for true proprioceptive content. Seeing your elbow bent or feeling pressure in the ventral skin and tension in the dorsal may provide the same propositional knowledge—“my elbow is bent”—but via a very different pathway to that of proprioception, which relies on a specific kind of stretch receptor in the muscle (Binder, Hirokawa, & Windhorst, 2009). This seems a clear case of Executive Attention without phenomenal experience of local proprioceptive content.

A similar argument might be made in the case of *circadian rhythm*. Researchers recently discovered a third type of receptor in the retina—in addition to rods and cones—that senses light but sends its signal to the hypothalamus, where it is used to regulate the body clock (Barinaga, 2002). This meets the criteria for Liberal Attention, but if one considers body clock information to influence one’s Executive Attention decisions—a controversial proposition, I grant—then this would be another case of an unconscious sensory perception that is Executively Attended. Similar cases can be made from other senses beyond the Aristotelian five, such as

¹⁴² This task especially relies on proprioception if one’s eyes are closed.

¹⁴³ Fulkerson (2014) discusses some complexities here.

¹⁴⁴ Armstrong (1995, p. 248) takes proprioception to be phenomenal perception. Dainton (2008, p. 209) says “This is not to say that proprioceptive awareness is wholly sensory, but it is to an extent.” I am arguing that it is at least partially *sensory*, but it is not at all *phenomenal sensation*.

vestibular balance sense, carbon dioxide blood concentration (Bogen, 2007, pp. 790–791),¹⁴⁵ and the ability to sense the salt concentration in the blood.¹⁴⁶

How has the case against $A \sim C$ been mounted? Some authors reject $A \sim C$ because of prior definitional (4.3.2.3) or theoretical (4.3.2.4) commitments. Among those who explicitly argue against empirical candidates for $A \sim C$, serious flaws are apparent. For example, Prinz (2011) argues against putative cases of $A \sim C$. Yet he himself accepts what he calls “unconscious perception” (pp. 2-3), which basically amounts—on his own description, and on any reasonable interpretation—to Executive Attention (selection, further processing, influence) without phenomenal consciousness. His reason for denying that this is $A \sim C$ is that the lack of consciousness in subliminal perception is due to “attention deficit” (p. 4). In other words, if the content is not phenomenal, it cannot have been attended, since attention is the “gatekeeper to working memory” (p. 9), and thus to consciousness and report.¹⁴⁷ As far as I can ascertain, a convincing case has yet to be made against the strongest kinds of candidates I described above as conclusive instances of $A \sim C$.

6.2.3 Interim Conclusion

I conclude from the foregoing that we may be quite confident that the Combination $A \sim C$ is instantiated in the human cognitive economy. That limits the available Scenarios to just two: $A \supset C$ or $A \cup C$. What will decide between the two is whether the Combination $C \sim A$ is instantiated in the human cognitive economy.

¹⁴⁵ Bogen thinks there is a phenomenal experience of high CO₂ blood concentration though I would argue it is not the high CO₂ itself, but the straining musculature, pain in the muscles etc., that is experienced.

¹⁴⁶ On the interconnectedness of different sensory modalities, see Bleicher (2012).

¹⁴⁷ I consider this kind of argument in detail in Chapters 7 and 8. Further persuasive critiques of Prinz’s argument may be found in Barrett (2014) and Taylor (2013a, 2013b).

6.3 Consciousness Without Attention

In this section, I address some preliminary issues relating to how we might go about identifying the presence of consciousness and the absence of attention, and the ways in which consciousness might be unstructured by attention. I then embark upon a methodical quest in search of empirical instances of $C \sim A$, structuring the search by classes of ways that the Combination might occur, given my definitions of Executive Attention and phenomenal consciousness. I consider candidates from five such classes in this chapter, most of which are *global* (6.2), and find none of them conclusive. The next two chapters are dedicated to detailed consideration of a putative candidate for *local* $C \sim A$ —the phenomenal content that overflows Executive Attention implemented in WM. By the end I hope to have settled upon $A \cup C$ as the most plausible Scenario and answer to Q.

6.3.1 Preliminaries

What would it mean for the content of cognition to be phenomenally conscious without being in any way Executively Attended? In exploring $A \sim C$ above, I discussed the difficulties in identifying the presence of attention and the absence of consciousness. Here we must consider the mirror image of those difficulties, namely, the difficulties in identifying the absence of attention and the presence of consciousness.

In searching for $C \sim A$, it will not be enough to establish phenomenal consciousness in the absence of the executive function of just one or a few of the Operational definitions of attention enumerated in my taxonomy of attention (3.3.4.2), but of all of them; nor of just one kind of attention (3.2.3), but of all of them. For example, consciousness without selection may still be consciousness with vigilance.

Consciousness without top-down attention may still be consciousness with bottom-up attention. In fact, what we seek is consciousness (local or global) in the absence of any implementation of any commensurate attentional strategies by the cognitive executive.

Identifying the presence of Executive Attention is in some ways easier than identifying its complete absence. Empirically, absence is also usually inferred from report by the subject, or by observing their physical behaviour, or by measuring physiological parameters such as the EEG, MEG, fMRI, etc., correlates of attention (or rather, the absence of such correlates, in our search). But this is notoriously difficult, and it may be one reason why researchers such as Koch and Tsuchiya (2007) resort instead to the somewhat less satisfactory “near-absence of attention” (4.2.4.2). A subject might not report their attending, might not exhibit any overt behaviour associated with attending, and might not show clear neural signatures of attending, yet still be attending (Fazekas & Nemeth, 2018).

Identifying the presence of consciousness also presents some major challenges. In cases where report is reliable, we can take that as a valid measure. But there are situations when we have good reason to question the reliability of report. I address some of these situations as they arise in the discussion below. Objective measures of phenomenal consciousness—the neural correlates of consciousness—are very controversial at present, as I observed above. I devote much discussion to this problem in Chapter 7.

In view of all these empirical challenges to identifying $C \sim A$, the discussion below must be taken as preliminary, and awaiting better methods of confirmation. It is nonetheless useful in its own right. A sound theoretical argument for a case of $C \sim A$, even in the absence of any conclusive empirical confirmation that such a case exists, still tells us something important about the relationship between attention and consciousness. At the least, it narrows the options by eliminating the Identity Hypothesis (4.4.1) by showing that there may be principled reasons why the two can come apart. It also points the way to where we should be looking for empirical

cases of $C \sim A$ and stimulates discussion about how the challenges above might be addressed in the different candidates under discussion.

Attention structures cognitions. In searching for $C \sim A$, we must be careful to remain focused on attentional strategies structuring (or not) phenomenal contents of cognitions, and not those structuring the processes that *produce* that content. As I have observed often, Liberal Attention indubitably contributes to any phenomenal content in the pre-phenomenal stages of its construction. What I will take to constitute genuine $C \sim A$ here is unified bound phenomenal content that becomes conscious “all at once” (Bayne, 2010, p. 238) that is not itself Executively Attended—that is, it is not *further* structured by Executive Attention.

It might be argued that merely being experienced itself constitutes a kind of attention-as-access. But I would argue (a) that this would be to fall back into a Phenomenal definition of attention, which I dismissed in 3.3.3, and (b) that consciousness need not be *reflective* in a higher-order way, but may be merely *reflexive* (2.4.1), therefore requiring no Executive Attention-as-access in order to be phenomenal.¹⁴⁸

The situated first-person perspective condition for consciousness can itself be taken as smuggling attention intrinsically into consciousness. In being situated, there is the inherent selection of taking one perspective to the exclusion of all other possible perspectives. Is not this selection and exclusion an implementation of two attentional strategies? And if so, is not attention thereby necessary for consciousness?

I would argue that being situated in this way is Liberal Attention, but not Executive Attention, unless we subscribe to a Phenomenal Definition of attention. Executive Attention is attentional strategies *implemented by the executive*. Unless that situated first-person perspective is held necessarily to arise only from the activity of the

¹⁴⁸ See also 8.6.2.2 and 9.5.2.1.

executive (which amounts to the Phenomenal Definition of attention), we cannot assume that Executive Attention is responsible for it. It may be capable of arising by means independent of the executive, and if so, then the fact that conscious experience has a selective and exclusive perspective in no way requires that it be Executively Attended.

6.3.2 Five Ways

I turn now from the conceptual analysis of $C \sim A$ to the quest for empirical instances of it. There have been many cases claimed to be $C \sim A$ in the literature. Some of these are cases where attention is not completely absent, but merely impaired, for example: “ADHD, depression, schizophrenia, bipolar disorder, post-traumatic stress disorder, and traumatic brain injury” (Rosenberg et al., 2017, p. 291),¹⁴⁹ and right parietal extinction and Balint’s Syndrome (Hardcastle, 1997, p. 59). Koch and Tsuchiya (2007) go further, and discuss a number of candidate cases where attention is not merely impaired, but “near-absent.” This list is: pop-out in search; iconic memory; gist; animal and gender detection in dual tasks; and partial reportability (p. 17, Table 1). Jennings (2015) considers three candidates: perceptual gist; imagistic consciousness (superblindsight, the colourblind colour expert); and phenomenal consciousness (as opposed to Block’s access consciousness). She goes on to propose her own candidate, “conscious entrainment.”

My quest for empirical cases of $C \sim A$ will range more broadly than those I have found in the literature, but not claim to subsume every case considered by others. Instead, I derive my list of candidates based upon the premise I set out above, that $C \sim A$ is

¹⁴⁹ These authors also discuss the possibility of using fMRI to distinguish or measure the attentional abilities of different individuals, as opposed to group-level data about attentional mechanisms in general.

phenomenal content that is unstructured by Executive Attention for further processing. Thus, I consider five broad classes of ways that content might be so unstructured, and search for empirical cases exemplifying each class. I do not claim these classes to be exhaustive—perhaps others might be thought of. And some of the candidates considered below might plausibly be interpreted in different ways and so fit under more than one class. But we only need one candidate to prove convincing to establish that $C \sim A$ is capable of instantiation in a cognitive economy. The five classes are: Pure Consciousness; global Unprocessed Content; Simplicity; Chaos; and Timing.

6.3.3 First: Pure Consciousness

“Some persons can voluntarily empty their minds and think of nothing” (James, 1890, p. 404).

The first candidate for $C \sim A$ is the possibility of global phenomenality without content, subjectivity without any object, or a consciousness that is not a “consciousness-of” anything. I will call this *Pure Consciousness*.¹⁵⁰ The concept of

¹⁵⁰ The idea I am considering here is closely related to other ideas in the literature, such as “C” (Bogen, 2007)—the core of experience, also called “subjectivity, sentience, awareness, consciousness itself, consciousness-as-such, consciousness per se, primary consciousness, or simply, by some earlier authors, consciousness” (p. 775), and *Minimal Phenomenal Selfhood* and *Minimal Phenomenal Experience* (Alcaraz, 2018). It may also be related in interesting ways to certain forms of panprotopsyism, given that the basic “units” of consciousness on these models may be too protean to bear any relationship to an object of consciousness. It is not related to “*minimal consciousness*” in the developmental sense (Zelazo, 2004), or to the *Minimally Conscious State* in the physiological sense (Giacino et al., 2002).

Pure Consciousness is quite controversial. It may even be an incoherent concept.¹⁵¹ But it is plausible if one accepts the distinction between phenomenality as such and the content of phenomenal experience (2.5). If such a thing is possible, then this kind of experience would have no content amenable to being structured, and therefore, no arena for attention to be implemented as a strategy for structuring it. It would be *C~A*. After briefly noting some unlikely candidates I consider four more interesting candidates: sensory deprivation; decortication; dreamless sleep; and meditation.

6.3.3.1 *Unlikely Candidates*

Consider the case of *Attention Deficit Hyperactivity Disorder* (ADHD) and related conditions. At first blush it seems the subject “loses attention,” but that is not, of course, what is really going on. ADHD is a failure, not of attention itself, but of *maintaining* attention on a single target over a prolonged period of time. The subject continually attends, but the focus of her attention jumps erratically from one target to another, making it difficult to complete set tasks. The same might be said of *mind-wandering* (Christoff et al., 2018; Irving & Thompson, 2017; Seli et al., 2018). One’s cogitative rambles do not lack structure, they lack a plan (Irving, 2016).

6.3.3.2 *Sensory Deprivation*

A more plausible candidate for global *C~A* is *sensory deprivation*: “to reduce the level of sensory stimulation to as low a degree as possible (e.g., darkness and silence)”

¹⁵¹ Kriegel (2007, pp. 130–131) expresses reservations as to whether the “method of sweeping imaginative subtraction” can really settle anything or be more than an intuition pump. I treat pure consciousness here as more than a thought experiment, exploring actual empirical candidates that involve phenomenal subtraction.

(Zubek, Hughes, & Shephard, 1971, p. 282).¹⁵² In the absence of any sensory input, perhaps as in a modern sensory deprivation chamber, there is no sensory input for attentional strategies to structure. But this too fails to constitute a case of $C \sim A$, for at least three reasons. First, even in the absence of sensory input from the outside world, the brain continues to generate its own “inner model.” Without the correction of constant sensory input, this model eventually begins to go awry, resulting in various kinds of aberrant thought patterns, including hallucinations, dreams, and body image changes, as well as emotions such as loneliness, tedium and anxiety (Zubek et al., 1971, p. 285 Table 1). This is *not* Pure Consciousness, it is consciousness with attentionally structured content that is less connected to the outside world.

Second, I know of no experimental paradigm that can prevent *interoceptive* sensation—sensory input from one’s own bodily organs—hunger, aching, etc. A high-level surgical or anaesthetic deafferentation might block interoceptive content, but there is a serious question about whether such a thing is possible (let alone ethical)—I suspect that any block effective enough to stop all interoception would also interfere with basic functions like breathing.

Third, attention-as-vigilance/monitoring is constantly active. Even in the case of prolonged sensory deprivation, one never loses the *ability* to sense input. So, the moment that new sensory input arises, the subject will immediately become aware of it, because her executive vigilance mechanisms are still functioning.

Sensory deprivation deprives experience of afferent perceptual content, but of course, the stream of consciousness includes more than that. There are memories and imaginings, thoughts and desires, to name a few. Can these types of phenomenal content be eliminated? Three possible candidates present themselves: decortication, sleep states, and meditation.

¹⁵² I consider perceptual deprivation below in the discussion of simplicity.

6.3.3.3 *Decortication*

Decortication is the absence of the forebrain. It occurs in at least two empirical situations: *anencephaly* (and related conditions), a congenital condition in which the forebrain simply does not develop in utero, and *decortication*, the surgical removal of large parts of the forebrain, sometimes used as a therapy of last resort for dangerous, intractable epilepsy. It is clear that in some patients, a surprising degree of ability can be preserved in the absence of large proportions of cortex (Merker, 2007; Swancer, 2017). These abilities almost certainly implement Executive Attention. What we are really looking for here is whether enough of the brain can be absent to eliminate all phenomenal *content*, yet maintain phenomenality as such. Unfortunately, we know too little about the neural correlates of consciousness, or the necessary and sufficient neural conditions for phenomenality to determine how much of the brain can be absent before either all phenomenal content or phenomenality itself is lost. Further, there are serious and likely insurmountable obstacles to identifying whether or not such patients are indeed phenomenally conscious (7.4), so this seems to be a currently unanswerable question.

6.3.3.4 *Dreamless Sleep*

When one sleeps and dreams, there is very plausibly *A&C* in the stream of dreaming experience. The brain undertakes a great deal of subconscious reorganisation and maintenance work in sleep (Nielsen et al., 2015),¹⁵³ some of which may constitute *A~C*. *Dreamless sleep* is generally taken to $\sim A\sim C$. But, some recent work (Siclari, LaRocque, Postle, & Tononi, 2013; E. Thompson, 2015a; Windt, Nielsen, & Thompson, 2016) suggests that one might be in some way phenomenally conscious in dreamless sleep, although the absence of dreaming, or perception, or motor

¹⁵³ See also the whole July 2015 issue of *Neurobiology of Learning and Memory* (122).

command suggests a consciousness with attenuated content, or perhaps even devoid of content completely. One of Windt et al's (2016) three categories of dreamless sleep experience is "'Selfless' States and Contentless Sleep Experiences" (p. 873, Box 1, see also p. 878), described as "a form of conscious awareness during sleep that lacks imagistic or propositional contents (E. Thompson, 2015a, 2015b; Windt, 2015)," and bearing a close affinity with the Indian philosophical idea of bare conscious awareness "without the subject-object structure of ordinary experience and the phenomenology of being a cognitive agent" (p. 878).¹⁵⁴

Windt et al., point out that experienced meditators who report "witnessing sleep"—a kind of lucid dreamless sleep, or meta-awareness of a sleep experience devoid of specific thought contents or imagery—exhibit the electrophysiological signatures of selective attention and memory (enhanced gamma-band activity). Non-meditators and inexperienced meditators neither report this kind of "witnessing sleep" nor exhibit the electrophysiological signature. This suggests that experienced meditators have learned to Executively Attend to a kind of pure phenomenal experience that is normally utterly unattended, and is normally, therefore, $C \sim A$.

Contentless phenomenal sleep is a plausible case of $C \sim A$, but there are some caveats. First, in all normal forms of sleep (excluding deep coma, anaesthesia, etc.) a type of Executive Attention persists, even in the absence of cognitive content—vigilance. A subject in any phase of sleep can be awoken by a loud enough noise. There is some debate as to the which comes first: does one attend because one first became conscious of the noise, or does one become conscious of the noise through attention-as-vigilance (De Brigard, 2010; Mole, 2008a)? But the vigilance itself counts as a form of continuously implemented Executive Attentional strategy, whether one is experiencing phenomenal yet contentless sleep or not. So dreamless sleep will not really work as case of global $C \sim A$. General anaesthetic induction and recovery give us good reason to think that global Executive Attention and phenomenal consciousness switch off and on together.

¹⁵⁴ See for example, *Maitri Upanishad*, 6:19.

Second, it is hard to find a characterisation of *completely* contentless dreamless sleep. Windt (2015) characterises dreamless sleep as being the pure phenomenal experience of time passing, or duration, without any further content. But of course, experiencing time qualifies as phenomenal content. Third, there again seems to be an insurmountable obstacle to confirming whether dreamless sleep is indeed phenomenal. We can only tell if a sleeper is indeed experiencing phenomenality, by them reporting it as such (as in the case of the experienced meditator). But to report it, she must learn to Executively Attend to it, rendering it no longer $C \sim A$. There is no dream content to be structured, it is true, but knowledge of the lack of content is itself attentively structured. I consider this epistemic conundrum in more detail in 7.4.

6.3.3.5 *Meditation*

Certain forms of Eastern *meditation* (Lutz, Dunne, & Davidson, 2007) are candidates for global $C \sim A$. These involve the idea of learning, through long practice, how to shut down all cognitive processes, leaving behind a residual of pure, contentless consciousness—“the continuity of consciousness through deep meditative states in which all conscious activity is said to have halted” (Chadha, 2015, p. 113).¹⁵⁵ Nash and Newberg (2013) identify six stages in meditation, five of which have Executive Attention at their heart. For example, in the Method stage, “methods have been defined as ‘a family of complex emotional and attentional regulatory training regimes developed for various ends, including the cultivation of well-being and

¹⁵⁵ In the context of meditation, Baars (2013) calls this “silent consciousness,” Taylor (2003, p. 331) calls it “‘stillness’ ... an extended content-free experience,” Nash and Newberg (2013, p. 14, endnote 1) call it “pure consciousness” or “emptiness” (citing F. Travis & Shear, 2010), and Vliengenthart (2011, pp. 156–157) calls it “absolute unitary being.” Pure Consciousness may also arise in certain kinds of mystical experience, not necessarily attained through eastern meditative practice (Forman, 1999).

emotional balance’ (Lutz, Slagter, Dunne, & Davidson, 2008, p. 163)” (p. 4). The sixth stage is the Enhanced Mental State (EMS), “the causal result of the successful application of the Method—an altered state of consciousness, commonly referred to as the meditative state” (p. 5). One form of this EMS is the Null State, “an enhanced empty state that is devoid of phenomenological content—a non-cognitive/non-affective state” (p. 6).¹⁵⁶

While this sounds very much like what I have called Pure Consciousness,¹⁵⁷ it actually fails to be case of *C~A* because it is not genuinely contentless. The self in this state does not disappear, but its *boundaries* seem to expand, merging the self with external reality—a much larger self, but a self nonetheless (Forman, 1999, p. 641). The situated first-person perspective remains, albeit with a different locus. Nash and Newberg speculate that this Null State of “non-self, or emptiness” may arise from the dampening of signals to and within the areas of the brain responsible for a subject’s sense of the *spatial location* of the self within the environment.¹⁵⁸ I argued above (6.3.1) that the presence of a situated perspective can arise through Liberal Attention, and does not in itself indicate the instantiation of Executive Attention. But if there is any sense of reflection upon that situated perspective, anything like the thought, “my self is enlarged,” or “I feel different,” or similar, that would indicate the kind of further processing that identifies an instantiation of Executive Attention. The descriptions I have read in the literature thus far do not make it clear whether such thoughts are completely banished in the Null State.

¹⁵⁶ Another tantalising description: “pure consciousness, true-Self, non-Self, NDA, absolute unitary being; and other terms such as Formless, Void, emptiness, and undifferentiated ‘beingness’ or ‘suchness’ ... Osho describes samadhi as ‘no object in the mind, no content, not meditating upon something, but dropping everything (so that) not even a ripple arises in the lake of your consciousness’ (Osho, 2003); and Sri Ramana Maharshi states that ‘samadhi is the state in which the unbroken experience of existence is attained by the still mind’ (Godman, 1985)” (Nash & Newberg, 2013, p. 8).

¹⁵⁷ We may assume for the sake of argument that the meditator can eliminate phenomenal interoceptions as well.

¹⁵⁸

We must beware of confusing “the perception of absence with the absence of perception” (Mole, 2008a, p. 98). The meditator in a Null State may not be attending to any particular content, but she is in a “hyper-attentive” state (Prinz, 2011, p. 183)—being aware that one is not perceiving anything is still a kind of Executive Attention (compare Ganzfelds, 6.3.5.3). Global *C~A* eludes us in even the emptiest of meditative states.

6.3.3.6 *Summary of Pure Consciousness*

To summarise the possibility of *C~A* due to an absence of phenomenal content altogether: sensory deprivation is only attenuation of phenomenal content, not total elimination. Decortication fails because we have no way of knowing whether phenomenality itself is snuffed out with the last bit of phenomenal content. Dreamless sleep seems plausible, although there is a major epistemic obstacle to overcome before we can be certain. Meditation fails because the inevitable awareness of the very absence of content is itself content. I conclude that none of these candidates is without serious flaws, and combining any two or more of these candidates serves only to multiply flaws rather than resolve them.

What if there is content (other than mere reflexivity), but it is unstructured, and therefore, unattended? I turn now to two ways in which this might be the case: global Unprocessed Content, where there is complex content that is not attentionally structured; and Simplicity, where there is so little content that attentional structuring is not possible.

6.3.4 Second: Global Unprocessed Content

Whereas Pure Consciousness is consciousness without content, *Unprocessed Content* is the phenomenal experience of perceptual or non-perceptual content that is not utilised in any form or fashion. One might purely experience content without any *further* access, selection, influence, or any kind of further attentional processing of that content.¹⁵⁹ I believe there is a “normal” variety of this, which I call First-Order Content, and will consider in detail in Chapter 7, in the context of a particular example of this class—the local perceptual content of phenomenal overflow. In this section, I consider two global candidates where none of the subject’s synchronic phenomenal content is further processed by Executive Attention (*C~A*)—an absence of what Schooler (2002) calls “meta-consciousness,” the second-order consciousness of one’s first-order conscious content. The two candidates I consider are mental blankness and stupor; and focal epileptic seizures.

6.3.4.1 *Mental Blankness and Stupor*

Mental blankness (hereafter, just “*blankness*”), total stupor,¹⁶⁰ vacancy, and oblivion are all terms that have been used for that state of mind where the subject experiences ongoing phenomenal perceptions, but is not thinking of *anything* at

¹⁵⁹ Some of the early stages of meditation—before one achieves Pure Consciousness—might fall under this category.

¹⁶⁰ As distinct from either coma or partial stupor (Joyston-Bechal, 1966, p. 969). A condition often associated with schizophrenia, severe depression, etc. For the development of the modern concept of stupor, see Berrios (1981).

all.¹⁶¹ Travis (1937) found that mental blankness¹⁶² was the only conscious state to show predominantly “large” EEG brain waves. This he interpreted as representing a “basic cortical equilibrium” (p. 307). The smaller, more “choppy” waves represent “a relatively high degree of specificity in psychic activity” (p. 309). This specificity is interestingly described as *focusing consciousness*, a kind of attention. For blankness to be a case of global C~A, we need to establish the absence of Executive Attention and the presence of phenomenal consciousness.

On the attention side, blankness seems plausibly to involve the subject’s disengagement of any Executive Attentional strategies from her synchronic phenomenal content. There is no selection or detection and recognition, in that no items within the phenomenal field are selected out for focal processing, and therefore nothing is actively recognised, neither is even the whole field recognised as such. The subject does not think to herself, “I am now blank.” She does not think anything at all. She merely experiences the state of being blank. There is no control, in that nothing is done with any of the phenomenal content—it is simply experienced. Neither is there coherence, in that no strategies are being implemented to make possibly incoherent content more tractable. So, apart from simple access in the sense of experiencing the content phenomenally in a first-order way, it appears that blankness is a strong contender for a case of global C~A.

But there are a number of problems with this account. First, about half of stupor patients have partial or total recall of things they experienced during episodes

¹⁶¹ This is different from a fugue state or certain kinds of amnesia, in which the subject pays attention throughout the fugue state, and acts in the world accordingly, but is unable afterwards to remember anything she did while in that state. Blankness is a kind of mirror image of the fugue state: there is little or no attention during the state, but that state is remembered afterwards. Neither is blankness the same thing as the default mode network, which “is active when individuals are engaged in internally focused tasks including autobiographical memory retrieval, envisioning the future, and conceiving the perspectives of others” (Buckner et al., 2008, p. 1).

¹⁶² Defined by Travis rather loosely as the subject answering a request to report what was going through their mind at a moment in time with answers like, “‘nothing,’ or ‘I’ve forgotten,’ or ‘nothing in particular’” (p. 305).

(Joyston-Bechal, 1966, p. 975, Table X). These experiences are often more than pure first-order perceptions, including complex if pathological cognitions (Berrios, 1981, p. 678). This suggests that although the patient's behaviour may indicate an absence of Executive Attention, that is not the case. There remain the proportion of cases where there is total amnesia of the experience, but here, even if Executive Attention is absent, it is impossible to know whether consciousness is present, which means this could well be a case of $\sim A \sim C$.

Second, subjects are responsive to the right kind of stimuli that can “shake” them out of the blank state. Executive Attention-as-vigilance remains intact to varying degrees in blankness. A more profound variety of stupor that might circumvent this concern is some forms of *catatonia* (Walther & Strik, 2016)—a clinical condition present in various psychiatric disorders—in which the patient is completely non-responsive, thus suggesting an absence even of attention-as-vigilance. However, responsiveness requires many stages of processing, including further processing of perceptual content, decision making, and motor command, not to mention motivation to respond (recall, some catatonic patients are deeply depressed). It only takes a failure of one of these steps to produce non-responsiveness, even if the remaining steps—all constituting active Executive Attention—are intact. It would be hasty, then, to infer $C \sim A$ from catatonic non-responsiveness. Perhaps neural studies of patients suffering from this condition might be able to shed light on the question, but I have not been able to find any such studies.¹⁶³

On the consciousness side, the partial or complete recall of some stupor patients (above) suggests that they at least were phenomenally conscious during the episode. But there remains a subset of patients with total amnesia of the episode. Here, it is

¹⁶³ Walther and Strik (2016, pp. 344–345) discuss the obstacles that have hitherto hindered such studies from being successful, before summarising those that have been carried out. Many of the studies they discuss are not relevant to the question at hand as subjects were not exhibiting stupor, although there was a tantalising finding of frontal and parietal hypoperfusion in patients with akinetic catatonia. The prefrontal cortex is often implicated in attention, so its hypoactivity may indicate a diminution or even lack of attention.

difficult to know whether this is an inability to *recall* actual phenomenal experience or a lack of phenomenal experience in the first place.¹⁶⁴ EEG and fMRI interpretation in stupor is complex and complicated by the influence of the underlying pathology (Brenner, 2005; Harrison & Connolly, 2013), and therefore unlikely to settle this question. Again, we are faced with the problem of relying on report (Executive Attention) to identify consciousness in the absence of Executive Attention (7.4). All these considerations make blankness in all its varieties unconvincing as cases of *C~A*.

6.3.4.2 *Focal Epileptic Seizures*

Epilepsy is a pathological condition that results from uncontrolled electrical activity in the brain. Whereas *generalised* epileptic seizures involve both cerebral hemispheres, *focal* seizures are unilateral and either discretely localised or more widely distributed (Berg et al., 2010). A recognised feature of some focal seizures is impairment of physiological consciousness,¹⁶⁵ which interestingly, is not a graded affair here, but bimodal (Cunningham et al., 2014)—patients tend either to be unconscious or lucidly conscious, with little gradation in between. Epileptologists have developed a bi-dimensional model of physiological consciousness in epilepsy with “level of consciousness” and “content of consciousness” dimensions (Monaco, Mula, & Cavanna, 2005), with psychometric instruments to evaluate the contents of consciousness in epileptic seizures (Alvarez-Silva, Alvarez-Rodriguez, & Cavanna, 2012, p. 512), and a third “self” dimension has also been proposed (Hanoğlu, Özkara, Yalçiner, Nani, & Cavanna, 2014). A distinction in the level of consciousness is also drawn by some between “non-consciousness,” “primary consciousness” (roughly,

¹⁶⁴ See Johanson et al., (2003, p. 284) for discussion of a similar problem in focal seizures.

¹⁶⁵ Focal seizures have also been called *dyscognitive* seizures (Blume et al., 2001). They should be distinguished from absence seizures, which tend last only a few seconds, and are more common in children (Panayiotopoulos, 2008).

first-order cognitions), and “reflective consciousness” (roughly, introspection) (Johanson, Revonsuo, Chaplin, & Wedlund, 2003, p. 280).

Some descriptions of certain kinds of focal seizures are plausible candidates for a kind of *C~A* similar to that of total stupor. The seizure patient is unresponsive suggesting a shutdown of Executive Attention, but physiologically conscious, suggesting they may also be phenomenally conscious (Monaco et al., 2005, p. 153).¹⁶⁶ By comparison, in *locked in syndrome* (Kondziella, 2017), the patient is richly conscious with an active executive, but efferent commands fail to produce any kind of physical action.¹⁶⁷ In focal seizures, it is possible that the executive itself is inoperative, thereby rendering Executive Attention absent. Given the depth of non-responsiveness, it seems that here, even attention-as-vigilance is inoperative.

The literature on this topic suffers greatly from an imprecision and inconsistency in the use of terms for attention and consciousness (Monaco et al., 2005, p. 150), and there is great diversity of symptomatology of focal seizures (Seneviratne, Woo, Boston, Cook, & D’souza, 2015, p. 591 Table 2). When an epileptologist speaks of “focal seizures with impairment of consciousness or awareness” (e.g., Berg et al., 2010, p. 260) it is difficult to be certain whether this means that phenomenal consciousness is impaired, or the usual Executive Attentional concomitants of conscious states are impaired. Confusingly, when Monaco et al., (p. 156) describe *limbic status epilepticus*¹⁶⁸ patients as having a high *level* of consciousness with no conscious *content*—which sounds like Pure Consciousness-style *C~A* (6.3.3)—they consider them to be cases of Searle’s (1992) philosophical zombies, which I would take to be *A~C*!

¹⁶⁶ They use the older terminology of *complex partial seizures*.

¹⁶⁷ A recent meta-analysis (Kondziella, Friberg, Frokjaer, Fabricius, & Møller, 2016) found that a significant proportion patients in a vegetative state exhibit evidence of physiological consciousness (14% by an active paradigm, 26% passive paradigm), and even more patients in a minimally conscious state (32% and 55% respectively).

¹⁶⁸ Different to focal seizures.

In descriptions by focal seizure patients (Johanson et al., 2003)—where the consciousness explored is “subjective experiences” rather than “objectively verifiable responsiveness” (p. 279)—the picture painted is as follows (pp. 280-284). Most patients experience a fall then rise in the *level* of consciousness over time. While some episodes involved reflective consciousness or primary consciousness for at least part of the episode, 65% of episodes involved a total absence of consciousness during the core event ($\sim A \sim C$). In many episodes, Executive Attention is not absent, but phenomenally experienced as being abnormally implemented ($A \& C$)—e.g., difficult to control or produces distorted content—and is inferred by the authors to be due to *abnormal* activity in the pre-frontal cortex, not *diminished* activity (p. 284).¹⁶⁹ It is difficult to extract a clear case for global $C \sim A$ from this picture.

Finally, we again stumble upon the epistemic obstacles related to amnesia and reliance on report. In focal seizures, amnesia may be anterograde (up to 5 minutes) and/or retrograde (up to 30 seconds) (Devinsky & Luciano, 1991), which Monaco et al., correctly point out makes the ascription of loss of consciousness during the seizure unreliable on that basis.

In summary, when the confusing terminology and subjective descriptions in the focal seizure literature are carefully analysed, we find evidence for $\sim A \sim C$ and $A \& C$ for at least some of the episodes, but no clear case of the $C \sim A$ we are looking for. Episodes of purely primary consciousness without reflective consciousness are described as implementing attention, albeit in abnormal ways. While this means that on the currently available evidence, focal seizures fail to establish $C \sim A$, there is great scope for further work to clarify the ambiguities described above and perhaps reverse this conclusion.

¹⁶⁹ Compare Mole’s (2014, p. 52) characterisation of Wu’s (2011, 2014) model of attention: “a subject is paying attention just in case she is navigating through the space of possible bodily and mental actions that are currently available to her, rather than drifting through that space at random.”

6.3.4.3 *Summary of Global Unprocessed Content*

Neither mental blankness / stupor nor focal epileptic seizures convincingly establish $C \sim A$. Putative cases are found either to involve Executive Attention of some kind or be epistemically opaque as to whether the subject is indeed phenomenally conscious.

6.3.5 Third: Simplicity

Another way that we might find $C \sim A$ is if the content of conscious experience were so utterly simple that it lacked sufficient content or complexity to be amenable to structuring. One can structure a Lego set in many interesting ways, but not a single Lego block. This situation might occur in at least three ways: inchoate perception; single object perception; and ganzfelds.

6.3.5.1 *Inchoate Perception*

An *inchoate perception* is one where the content is so sparse as to make it difficult or impossible to determine with any confidence what the exact nature of the stimulus is. Let us say for illustrative purposes that a patch of skin on the upper arm of 50mm diameter has only one tactile mechanoreceptor. Without looking, a subject cannot distinguish where in the receptive field she is touched by a blunt needle-end, or even tell the difference between being touched by one needle or two needles simultaneously 30mm apart, since the single mechanoreceptor fires identically in both cases. This inchoateness contrasts with the very clear and precise high-resolution tactile sensation at, say, the very sensitive, mechanoreceptor-dense finger-tips. No amount of attention can refine inchoate perception or improve its

resolution, since the data just isn't there. Is this a case of local $C\sim A$ in inchoate tactile sensation?

In response, it could be argued that this is not a case of the *absence* of attention, but of a *failure* of attention. When the subject strains to make out whether it is one or two needles, the subject is certainly Executively Attending, implementing selection, increased access, and spotlighting of attention onto the sensations coming from that dermal region. The only problem here is that all this attention fails to produce the desired result: discrimination between the two needle-pricks. It is not the case that one is conscious of the Targeted two needle-pricks, but unable to attend to them as two needle-pricks. Rather, one is neither conscious of, nor able to attend to, two needle-pricks. So, this is not a case of $C\sim A$, but of local $\sim A\sim C$. I consider an inchoateness account of $C\sim A$ in 8.6.3.

6.3.5.2 *Single Object Perception*

The second candidate for $C\sim A$ due to simplicity is the local¹⁷⁰ case of visual *single object perception* (Koch & Tsuchiya, 2007). If one's visual field contains only one object, then there is no scope for selection or exclusion, since there are no other objects to select among or exclude. But this is not really an absence of attention. One may, for example, be exercising spatial selection, attending to the locus of the single object in space, rather than to the other loci in the empty space around it. Further, any visual object has features: edges, dimensions, etc. The subject can attend to some features to the exclusion of others. Even a single dot at the boundary of visual resolution still presents a contrast against the background that can be the Target of feature attention. This too is not $C\sim A$.

¹⁷⁰ Not global, since attention in non-visual modalities remains unaffected.

6.3.5.3 *Ganzfelds and Perceptual Deprivation*

The third candidate for $C \sim A$ from simplicity is the local case of visual experience of a *ganzfeld*.

“Consider, first, an argument from Cristof Koch (personal communication). Imagine you are looking at an equally luminous, equally saturated wall of color (a “ganzfeld”). Your entire visual field is taken up by the color, and there is no variation in it, and no objects to focus on. In this case, there is no need to allocate attention, because attention is a selective capacity, and there is nothing to select. So, Koch reasons, under such conditions, attention is not engaged. Yet, it is obvious we would experience the color” (Prinz, 2011, p. 196).¹⁷¹

Prinz responds (correctly, I think) that (a) one may allocate spatial attention to different parts of the visual field, even if there is no object or feature attention; and (b) it is just as much attention to attend to the field as a whole as it is to attend to one object in a visual field. For Prinz, attention is not solely that act of selection, but the bringing of information into WM—i.e., access is also involved. Since the perception of the single bland wall of colour is indeed brought into WM, then Executive Attention is implemented here, even in the absence of object selection or competition. All this highlights the benefit of having a broad view of attention like the one I developed in Chapter 3—Koch’s intuition in this exchange derives from too narrow a definition of attention.

¹⁷¹ Tononi and Laureys (2009, p. 379) make the same argument.

Ganzfelds are a visual variety of *perceptual deprivation* (I discussed its experimental sibling, sensory deprivation, above in 6.3.3.2).

“the objective of the PD [perceptual deprivation] procedure is to provide homogeneous and unpatterned stimulation (e.g., diffuse light and white noise) while maintaining the level of sensory input near normal” (Zubek et al., 1971, p. 282).

The ability of perceptual deprivation to produce hallucinations has been well known for some time and continues to be an object of interest to researchers.

Electrophysiological studies have implicated attentional processes in the production of these ganzfeld hallucinations, with the EEG signature of attentional activity, low-frequency alpha waves, being recorded in the brains of subjects experiencing them (Wackermann, Pütz, & Allefeld, 2008). This is clear evidence of Executive Attention implemented in perceptual deprivation—albeit, producing abnormal content—therefore making them cases of *A&C*. Perceptual Deprivation, too, fails to establish *C~A*.

6.3.5.4 *Summary of Simplicity*

In inchoate perception there is a commensurate attenuation of both Executive Attention and phenomenal consciousness. In single object perception, there is at least spatial Executive Attention, even if it has only a single Target. And in Ganzfelds and Perceptual Deprivation, there is at least spatial Executive Attention, even if its only Target is a blank field or a multimodal perceptual field without specific content other than its emptiness (which constitutes attendable content of a sort). None of these cases conclusively establishes *C~A*.

6.3.6 Fourth: Chaos

“[Attention] implies withdrawal from some things in order to deal effectively with others, and is a condition which has a real opposite in the confused, dazed, scatter brained state which in French is called *distraction*, and *Zerstreutheit* in German” (James, 1890, p. 404).

In some focal seizures (6.3.4.2) attention malfunctioned, producing abnormal content. Might this kind of *chaotic* phenomenal content be taken as the “real opposite” of attention, as James suggests? Pure Consciousness and Simplicity putatively lacked attention because there is little or no content to be attended. Unprocessed Content putatively lacked attention because content was present, but not processed in any higher-order sort of way.¹⁷² But what if first-order content is present, and further processed, but this higher-order processing lacks the implementation of any Executive Attentional strategies? The result would be *Chaotic* phenomenal content, a putative case of local $C \sim A$.

Note that attentional strategies are not solely responsible for the ordering of content. For example, binding is crucial to ordering the content of phenomenal experience, but as I argued in 3.4.2, binding itself is not an *attentional* strategy. So, chaos could theoretically arise from a failure to bind, even with Executive Attention active and intact. The unlikely candidates considered in 6.3.3.1 above could plausibly be construed as Chaotic cases, but would fail here for much the same

¹⁷² The result of unprocessed content is not chaos, but stillness. One perceives the world as it is, and that’s it, full stop. Malfunctioning Executive Attention, on the other hand, involves not stillness, but chaos. Cases of Unprocessed Content would be quite calm to experience, while cases of chaos would seem quite disturbing to the subject.

reasons discussed above. In this section I consider two stronger candidates for $C \sim A$ —psychedelic states and simultanagnosia.

6.3.6.1 *Psychedelic States*

It is uncontroversial (I think) that subjects undergoing *psychedelic states*—induced by psychoactive drugs, for example—have a phenomenal experience of the cognitive content of that state. But such content is often described as being different from our normal waking states, in that it is disjointed, fragmented, or chaotic. Might this chaos be an indication of an absence of Executive Attentional strategies?

Letheby and Gerrans (2017) recently described an interesting model of ego-dissolution in the psychedelic state. On this model, the self is “not an actual entity or an object of perception, interoception, or introspection” (p. 2), but “an entity, substance, or bare particular that instantiates properties” (p. 1). Without going into the subtleties of this definition of a self which are immaterial to my argument here, the self is roughly described as a *model* of a cross-modal unitary ego, constantly being generated by the brain in order to predict the flow of experience. Attention, on this view, serves to make ideas “sticky” (i.e., increased access, selection) to this self. “At all levels, salience is attributed, attention directed, and information integrated in accordance with the relevance of information to the organism’s goals” (p. 9). The effect of the psychedelic drug is to inhibit certain attentional processes, thus—among other things—disrupting the “normal” binding process, disorganising the predictive model, and producing a diminished sense of self, one that differs from our normal experience by being less unitary or coherent.

Is consciousness present in a psychedelic state? Letheby and Gerrans make an important comment: “even in florid psychedelic experience the self-model is never entirely destroyed” (p. 2). That is, the conscious experience of being a self never disappears completely, but rather, its content is modified such that the connection of the present experience to the background content of selfhood is more tenuous.

Indeed, their view blurs the phenomenality-content distinction for which I argued in 2.5.

“The subsequent diminution in the sense of solid selfhood shows subjects that this sense is ultimately just one more conscious experience, rather than a transcendental precondition of all such experiences” (p. 9).

Thus, that component of the content of consciousness that is the self-model is not eliminated in the psychedelic state. Indeed, it can seem “both more intense and less personal, and salience, affective feeling, and motivation become detached from personal goals and history” (p. 6). It is the coherence of the content “surrounding” that self-model that is disrupted.

“Psychedelic subjects often report that their sense of being a self, or ‘I’, distinct from the rest of the world ‘out there’, is weakened, altered, or abolished during the intoxication ... Various authors have suggested that some vestiges of self-awareness are preserved in most, if not all, psychedelic experiences (Pahnke, 1969; Shanon, 2002), which helps explain the puzzling fact that autobiographical memories can apparently be formed of these putatively selfless episodes (Metzinger, 2005)” (p. 6).

On this description, it seems that what is dissolved is the sense of *separation* between the self and the rest of the world—a blurring of the boundary of the self,

while the sense of phenomenal selfhood itself—of being “me”—remains intact.¹⁷³ Further, while the content of conscious experience is severely disrupted in various ways, the self-model persists, even if in an impoverished-content state.¹⁷⁴ Much like meditation above, on my definition of consciousness, all four core characteristics remain intact, even if the locus of the situated perspective is somehow expanded, whether spatially or conceptually.

So, consciousness is preserved in the psychedelic state, but is there an absence of Executive Attention? That is hinted at, of course, by the failure of self-binding, which on their view, depends on attention. But there seem to be good reasons to think that attention is never completely absent in the psychedelic state. For one thing, the ability to form autobiographical memories mentioned in the quotation just above suggests that the content of psychedelic experience is not only accessible to, but actually accessed by WM, and convertible into verbal report. What is more, Letheby and Gerrans observe that

“Subjects often find their attention drawn to stimuli which they normally would not notice; as Watts puts it, psychedelics ‘make the spotlight of consciousness a floodlight which ... brings to light unsuspected details—details normally ignored because of their lack of significance’ (1964).¹⁷⁵ Attention is no longer guided exclusively by adaptive and egocentric goals and agendas; salience attribution is no longer bound to personal concern” (p. 6).

¹⁷³ “Ego dissolution experiences often occur in the context of mystical states in which the ordinary sense of self is replaced by a sense of union with an ultimate reality underlying all of manifest existence—the famous ‘cosmic consciousness’ experience” (p. 6).

¹⁷⁴ “ego dissolution is not an all-or-nothing affair. Different aspects of self-awareness may be more or less disrupted in different ways on psychedelics” (p. 6).

¹⁷⁵ The original in Letheby and Gerrans seems to mis-cite Watts as 2002 rather than 1964. Or perhaps they mis-cite 2002 as Watts instead of Shanon. Unfortunately, they don’t provide a page number.

It is not, then, that attention is absent in psychedelic states—it just operates differently. The goals subjects strive to fulfil are those of a different and larger self.¹⁷⁶ But the strategies of Executive Attention—a spotlight of focus, increased processing, alerting, salience—continue to be implemented. The chaos of psychedelic states, then, is not a failure of attention altogether, but a reordering of how attention operates. We will not, it seems, find *C~A* here, but *A&C* with unusual content, and where the situatedness of phenomenal selfhood is preserved but expanded, with blurred borders with the rest of the world.

6.3.6.2 *Simultanagnosia*

Another plausible candidate for *C~A* due to chaotic content is *simultanagnosia*. A recent review summarised the condition thus:

“Simultanagnosia is a disorder of visual attention that leaves a patient’s world unglued: scenes and objects are perceived in a piecemeal manner. It is generally agreed that simultanagnosia is related to an impairment of attention, but it is unclear whether this impairment is object- or space-based in nature ... These patients can see only one object at a time and sometimes only pieces of objects, unaware that they are locked on just one component of a larger form. What they can see cannot be located in space, likely because they see nothing else that can provide a reference point to

¹⁷⁶ “These results strongly imply that the subjective effects of psychedelic drugs are caused by decreased activity and connectivity in the brain’s key connector hubs, enabling a state of unconstrained cognition” (Carhart-Harris et al., 2012, p. 2138).

situate objects in the world” (Dalrymple, Barton, & Kingstone, 2013, p. 1).

Simultanagnosia, then, is a failure of the attentional strategies that normally underlie the formation and perception of a holistic synchronic experience containing multiple bound contents. Yet clearly, there is phenomenal experience of this chaotic content. Might this constitute a case of $C\sim A$?

There are at least two ways this might be $C\sim A$. The first is that the whole content of experience is, for the simultanagnosiac, globally unstructured by attention. But this is quickly discounted when we consider that the patient is indeed focusing attention upon the minimal content available to her (limited by the neurological deficit), and quite capable of attending to content in other modalities, so this is not a case of global $C\sim A$. It may also be a failure of binding rather than Executive Attention as such, which, as I argued in 3.4.2, are distinct cognitive strategies.

The second way one might find $C\sim A$ in simultanagnosia is to argue that the objects/space in the periphery is consciously experienced, but not Executively Attended in any meaningful way. Now there is some question as to whether this peripheral experience is phenomenally experienced at all. But even if it is, this becomes an argument for local $C\sim A$ —a kind of phenomenal overflow, which I consider in Chapters 7 and 8.

6.3.6.3 *Summary of Chaos*

In summary, then, neither of the putative cases of $C\sim A$ due to chaos stands up to scrutiny. In psychedelic states, Executive Attention is implemented but produces unusual Consequences, built around an expanded self. And simultanagnosia fails to be global $C\sim A$ because some attenuated content is normally Executively Attended,

while the content that is not needed to be considered as a special case of phenomenal overflow, which I consider in the next chapter.

6.3.7 Fifth: Timing

Obviously, attentional strategies take time to be implemented in cognitive processes. One indication that Executive Attention might be absent could be in situations where there is not enough time for its strategies to be implemented. If there is, however, enough time for conscious experience of that content to occur, we would have a case of $C \sim A$. Here I consider two candidates for this kind of $C \sim A$: gist perception; and the sequence of alerting.

6.3.7.1 *Gist*

“In a mere 30ms presentation time, the gist of a scene can be apprehended. This is insufficient time for top-down attention to play much of a role” (Koch & Tsuchiya, 2007, p. 18).

This makes rapid conscious *gist perception* a plausible candidate for local $C \sim A$. The key here is that the 30ms is only the duration for which the subject is *exposed* to the stimulus. While it may only take 30ms for the sensory data to register upon the retina, that content may be held in iconic memory for significantly longer, allowing plenty of time for Executive Attention to process it enough to arrive at a gist perception. The consciousness of that gist itself almost certainly also does not come about at the exact moment the 30ms exposure concludes, but sometime later.

Jennings (2015, pp. 278–281) makes a strong case for gist occurring without attention, but not *conscious* gist. This is not a case of $C \sim A$.

6.3.7.2 *Sequence of Alerting*

The temporal relationship between attention and consciousness presents another difficult conundrum. Consider an interesting case due to Mole (2008a, p. 90): when a baby’s cry awakens his sleeping mother, which comes first—was the mother first conscious of the cry, and therefore had her attention drawn to it; or was did the mother become conscious of the cry because her attention was first drawn to it?¹⁷⁷ If the former is correct—consciousness arises first and draws attention—then there is a brief window of time in which the mother is conscious of the cry, but not yet attending to it: local $C \sim A$.

Mole presents an interesting discussion of this conundrum, but on my definitions of Executive Attention, consciousness, and WM, we can formulate a relatively straightforward solution. The important propositions here are that vigilance and monitoring are attentional strategies; that implicit WM can operate without the benefit of consciousness; and that attentional strategies implemented by the executive of WM constitute Executive Attention. Given those propositions, it seems clear that Executive Attention implements vigilance and monitoring in the absence of phenomenal consciousness, and even in the severely attenuated physiological consciousness of sleep. Unconscious processes involving the normal auditory apparatus of WM first process the auditory stimulus, compare it to known sounds, identify it as the baby’s cry, and therefore “raise the alert,” which then activates the awakening of physiological consciousness bringing with it—and drawing the focus of attention to—the phenomenal experience of the baby’s cry. This goes against the intuitions of Mole’s sample of mothers (p. 91) but they don’t have the benefit of our

¹⁷⁷ Schwitzgebel (2007, p. 13) raises a related case of tactile experience of one’s foot in one’s left shoe.

definitions! This is not a case of $C \sim A$ —if anything, there is a brief temporal window in which Executive Attention is processing the cry before the mother becomes conscious of it— $A \sim C$.

6.4 Chapter Summary

In this chapter I first considered the possibility of $A \sim C$ and found that this Combination is not only conceptually cogent, but well supported empirically. This left us with two live Scenarios: $A \supset C$ and $A \cup C$, and the Combination $C \sim A$ as the key to choosing between them. I therefore devoted the balance of this chapter to considering it. Once again, $C \sim A$ is conceptually cogent, but is it empirically supported? Of the five putative empirical classes of cases, encompassing fifteen specific conditions considered in detail, none have so far proved to be incontrovertible cases of $C \sim A$. But there remains one more case that is the most promising of all those I consider: phenomenal overflow. It is to a detailed analysis of that case that I turn in the next two chapters, since the answer to Q would seem to hang on it.

7 Phenomenal Overflow

7.1 The Enigma of Phenomenal Overflow

The last of the putative cases of $C \sim A$ I consider is that of *phenomenal overflow* (henceforth, simply “overflow”), most prominently posed by Ned Block.¹⁷⁸ This is the idea that we consciously experience more than we can attend to, access, further process, or report (broadly construed, see 7.2.3). The content of our phenomenal experience at any given moment in time is therefore much richer than that to which we can Executively Attend. If overflow obtains, it is indubitably a case of local $C \sim A$, since the overflowing content—that subset of phenomenal content that is not Executively Attended—instantiates the Combination $C \sim A$.

At least partially in response, some have argued that what Block takes to be rich overflowing phenomenal content is in fact much sparser (e.g., Cohen et al., 2016b).

¹⁷⁸ Block (1995, 2001, 2005, 2007, 2011a, 2013b, 2013a, 2014b, 2014a, 2018).

The brain employs strategies such as ensemble statistics (Alvarez, 2011; Whitney, Haberman, & Sweeny, 2014) to create rough, sparse averages of the content peripheral (not necessarily spatially peripheral) to the focus of attention while focusing its processing resources on the content in the focus of attention to produce genuinely rich content there. However, since by definition we are not attending carefully to the content peripheral to attention, we never realise that it is in fact sparse, instead assuming that it must be as rich as the content in the focus of attention—an “Illusion of Richness”¹⁷⁹ (8.6.1)

There is both interest and plausibility to both accounts (and I consider some others in 8.6), and much ink has been spilt arguing over the matter. In this treatise, I essentially argue that while it is not yet possible to establish the fact of phenomenal overflow beyond doubt, recent attempts to discredit it—along the lines of Cohen et al., and others—have not been successful, and I explore some reasons for that failure. In fact, I argue that the overflow account offers the best explanation of the evidence currently available, especially in one particular case: foveal vision.

After some pertinent clarifications and definitions, I introduce Florence and her forest, and Ben and his hen as illustrative cases that highlight an important distinction that nuances Cohen et al.’s position: full-field vision must be treated differently to foveal vision. This brings us to the question of just how we can know, in the foveal case, whether all the content that is processed through the retina is experienced phenomenally, or whether only some of it is. I present three arguments for keeping an open mind. First, there are indisputable cases of diachronic phenomenal overflow, and what we take to be synchronic experience may just be a very short-duration case of diachronic experience. If so, then overflow becomes more plausible. Second, there is an incorrigible immediacy about certain kinds of very rich experience that simply cannot be extinguished or brought into doubt, no

¹⁷⁹ This is closely related to the “Refrigerator Light Illusion” on which one thinks the whole visual field is rich because whenever one turns one’s attention to any part of it, one finds that part to be rich (Schwitzgebel, 2007, p. 13). Other factors may be at work in addition to ensemble statistics, such as inflation (Odegaard et al., 2018).

matter how the subject introspects it. And third, an epistemic impasse inherent in empirical methods for finding out whether phenomenal overflow occurs suggests that not only have they failed thus far to settle the question, but that they may never be capable of doing so.

What I am attempting here is not a knock down deductive argument, but an abductive one: given the material discussed below, an open mind on phenomenal overflow is still in order, and it may even be the case that phenomenal overflow is a better explanation of the evidence than the illusion of richness or other overflow-sceptical accounts.

7.2 Some More Foundations

A little further groundwork is necessary, in addition to that laid down in Part I. In this section I distinguish between three kinds of cognitive content that are particularly germane to the overflow question, draw a connection between Block's "access consciousness" and my Executive Attention, and stipulate what I mean by "report," a concept that is also central to my analysis. I make no claims to mine being the only or even the "correct" definitions—I merely stipulate the particular (well-grounded, I hope) concepts that I think are interesting in relation to the question of overflow.

7.2.1 Three Types of Phenomenal Content

The key to addressing the overflow question, I believe, is to see that not all phenomenal content is equally structured for further processing by Executive Attention. Like most modern diets, it consists of some content that is very highly processed, some that is moderately processed, and some that is in its raw state. If this hypothesis can be established, phenomenal overflow becomes not only plausible, but highly likely. The relatively “raw” content—content that is not processed in any kind of higher-order way—is where we will find overflow, *if at least some of the unprocessed content is phenomenal content*. Much of this chapter will be devoted to the question of whether or not the “raw” content is phenomenal, and whether we can know it to be so.¹⁸⁰ In this section, I develop a distinction between three types of phenomenal content based on whether they undergo high-order processing, and what kind of processing they undergo.

One way of depicting this differential processing of content is to draw a distinction between three types of content. Consider this passage:

“The claim is not that we are unaware of our own conscious beliefs and experiences (or unaware that we have them). It is, instead, that *our being aware of them, or that we have them*, is not what makes them conscious. What make them conscious is the way *they make us conscious of something else*—the world we live in and (in proprioception) the condition of our own bodies” (Dretske, 1993, p. 281, my italics).

¹⁸⁰ Lest it be said I am begging the question here—I am not starting by claiming that there is such a thing as pure consciousness or phenomenal content that is not further processed in any way (indeed, see 6.3.3). That is a matter for empirical investigation. Rather, I start by claiming that phenomenal content that is not processed in any higher-order way is *conceptually* possible, and then proceed to seek empirical evidence that it in fact occurs.

For Dretske, there is conscious content that (a) makes “us conscious of something else—the world,” but this content does not rely for its phenomenality on things like (b) “our being aware of them” or (c) our being aware “that we have them.” We can formulate this distinction as one between three kinds of phenomenal content:¹⁸¹

First-Order Content (FOC) is that subset of phenomenal content that is not processed in *any* higher-order way. For example, to simply see a leaf that happens to lie in my direct foveal vision while my mind is fully engrossed in solving a complex mathematical problem is visual FOC. Visual iconic memory is an example of a cognitive system that may often give rise to FOC (Prinz, 2011, pp. 180–181).¹⁸²

Experience-of is *perceptual* content produced by a certain kind of higher-order processing of FOC. For example, to attend closely to the texture of the leaves I see, or to be aware of their particular shade of green are examples of Experience-of higher-order processing of that FOC. Experience-of is purely perceptual, on my definition. Experience-of content arises from higher-order processing implemented in a global workspace (Baars, 1988), or in the central executive–episodic buffer axis of multicomponent WM (Baddeley, 2003, p. 836). A possible example of relatively pure Experience-of might be generating visual imagery (e.g., of a hen) with one’s eyes closed (Schwitzgebel, 2011, p. 36).

¹⁸¹ Each of the three has its non-phenomenal counterpart.

¹⁸² I use the term “iconic memory” (Neisser, 1967) without any commitments as to the nature of its contents, representational or otherwise. For a discussion of the related concept of Fragile Visual Short-Term Memory, see Barrett (2014, p. 23).

Seeming-that is *conceptual* content produced by a different kind of higher-order processing of FOC (or of Experience-of) that can be characterised as demonstrative or propositional. It is that part of conscious content that is *about* what I see, whether about the FOC, or about other Experience-of type detail. Knowing that what I am looking at is a forest would be Seeming-that. Gist perception is an example of a cognitive process producing Seeming-that (Mack & Clarke, 2012).

Dretske (1993, pp. 276–277) (partially) makes the distinction clearer in his discussion of Rock’s wiggles. Subjects looking at two near-identical but slightly different cloud line-drawings often failed to notice the area where some wiggles were missing in one drawing but not the other. He describes this as subjects being “thing-aware”—they saw the different wiggles—but not being “fact-aware”—they were unaware of the fact that the wiggles were different. On my threefold distinction, there is a bare perception of the wiggles, a visual image formed in the mind—which is FOC—utterly bereft of identifications, comparisons, judgements, etc. If the subject is asked as to whether the clouds differ from each other, she implements higher-order processing of the visual image, Executively Attending to various features and locations, to produce a more refined Experience-of the clouds, perhaps including a perception of the wiggles that differ that groups them together for comparison. On this basis, she is able to come to conclusions or judgements about the fact that the wiggles are cloud-shaped, and that they are slightly different—Seemings-that.

7.2.1.1 *Similar and Supporting Distinctions*

These distinctions are surprisingly common intuitions that go back a long way. They are inherent in the reflexive/reflective distinction in Indian philosophy (2.4.1). First-order cognitions (e.g., perceiving a blue sky) can themselves become the objects of

second-order cognitions (e.g., I am perceiving a blue sky now, rather than yesterday). The second-order cognition is *reflective*, in that it is a cognition *about* another cognition within the same subject. However, inherent in the first-order cognition is that “I” am the subject, experiencing this cognition. This is *reflexivity*, and differs from reflectivity in that it is inherent in the very act of consciously perceiving that there be a subject doing the perceiving. This first-order, ineliminable subjectivity is reflexivity, whereas the optional (in that it doesn’t always happen) second-order consideration of the first-order perception is *reflectivity*.

According to Wider (1997), Descartes draws a distinction between *pre-reflective* and *reflective* consciousness, as does Sartre (pp. 11-14). Kant spoke of the idea of *transcendental self-consciousness* where no reflection is involved (pp. 35-38), and Locke spoke of degrees of self-consciousness, distinguishing, for example, between a thought and a perception of that thought (pp. 16-17).

More recently, others have drawn distinctions very similar to all or part of my three types of content (Bayne, 2010, p. 78; Bayne & Montague, 2011, pp. 9–10; Campana & Tallon-Baudry, 2013; Coltheart, 1980; Gallagher, 2004, p. 90; Gross, 2018, p. 2; Horgan, 2012, p. 408; Lamme, 2003, p. 14, 2010, p. 208, 2018, pp. 3–4; May, 2004; Mogensen & Overgaard, 2018, p. 6; Mole, 2011b, pp. 72–73; Naccache, 2018, pp. 2, 5; Schooler, 2002; Stoljar, 2016, pp. 1193–1195; Wolfe, 1994, 1999a, pp. 3–5).¹⁸³ The similarities and differences lie beyond the scope or the requirements of this treatise.

I take it then, that this hierarchical structure of cognitive processing is so widespread as to be largely uncontroversial. Few would disagree that cognitive content is processed in many steps, often (if not always) involving Liberal Attention. The controversial question here is whether there is a stage at which rich content has

¹⁸³ Compare also the auto-noetic/noetic/anoetic and the remember/know/confidence distinctions (H. Roediger, Rajaram, & Geraci, 2007; Tulving, 1985), the non-consciousness/primary consciousness/reflective consciousness distinction in epilepsy (Johanson et al., 2003, p. 280), and the distinction between low-level properties of experience (colours, shapes, etc.) and higher-level properties (potential, causal relations, etc.) (Masrour, 2011, p. 366).

become phenomenal but at least some of that phenomenal FOC is never further processed in any way employing Executive Attention—i.e., not processed by the cognitive executive. It is worth elaborating and elucidating the three types of phenomenal consciousness a little further in preparation for arguments below that shall build upon them.

7.2.1.2 *FOC*

Cohen and Dennett (2011, p. 362) ask, “What does it mean to have a conscious experience that you yourself do not realize you are having?” Phenomenal FOC is precisely that possibility (although a lot of the time, it is accompanied by Experiences-of and Seeming-that). That FOC can be phenomenal is guaranteed by its ability (in theory at least) to satisfy my four characteristics of consciousness (2.4). It is utterly coherent to say that I see wiggles with reflexive phenomenality—it being inherent in the experience that there is an “I”—a situated first-personal subject who is doing the seeing—without reflecting upon such seeing in any higher order way, and therefore without “realising” (it Seeming-that) I am seeing the wiggles.¹⁸⁴

To be clear: FOC is the result of processing, just not *higher-order* processing as I have defined it here (Executive Attention in WM). There is significant processing taking place at the retina and post-retinally in the early visual centres of the brain, and probably even beyond. This processing involves Liberal Attention. But the overflow question (as I am framing it) is not concerned with Liberal Attention, but with Executive Attention. FOC is content that is not (yet, at least, or perhaps, ever) Executively Attended. I allow that the extent of Executive Attentional processing

¹⁸⁴ This assertion is neutral with regard to what we might think subjects are, or even whether subjects are real entities at all. For a variety of views on this question, see Gallagher and Shear (1999) and Gallagher (2011).

may range from prodigious down to near-absence, but for the purposes of Q it would still be regarded as Executive Attention regardless of its “quantity.”¹⁸⁵

Also to be clear, Experience-of and Seeming-that by no means exhaust higher-order processing. Remembering the name of something you saw, calculating the sum of two numbers on the page in front of you, feeling sad on seeing a picture of a departed loved one—all these are examples of higher-order processing that are in themselves quite plausibly neither Experience-of or Seeming-that as such, yet they certainly involve some kind of higher-order processing (and therefore, Executive Attention).

7.2.1.3 *Higher-Order Processing: Experience-of and Seeming-that*

I have, however, focused on phenomenal content arising from two particular varieties of higher-order processing, because it is these that are most relevant to the overflow question. I noted earlier a difference between propositional *confidence* (3.3.3.2) and perceptual *clarity* or *vividness* (3.3.3.3). Bayne also describes this distinction between *conceptual* and *perceptual* higher-order content.¹⁸⁶

“This is perhaps most obvious with respect to change blindness. It is one thing to be aware *of* features in a scene that have changed, and it is another to be aware *that* they have changed” (Bayne, 2010, p. 78, italics in the original).

¹⁸⁵ Koivisto et al., (2009) explore the complex electrophysiological relationship between attention and FOC.

¹⁸⁶ Siewert (2012) similarly argues for both “sensing/thinking” and “first-order/higher order” distinctions.

Being aware *of* features is my Experience-of while being aware *that* they have changed is my Seeming-that. In change blindness paradigms, it is possible for subjects to have none, either, or both. Thus, we might think of FOC as being reflexive but not reflective, and perceptual; Experience-of as being both reflexive and reflective, and perceptual; while Seeming-that is both reflexive and reflective; and conceptual. In the realm of the phenomenal content retrieved from memory, Experience-of may be taken to be that which results from non-declarative memory retrieval while Seeming-that arises from declarative memory retrieval (see Squire, 2004, 2009).¹⁸⁷

The conceptual/perceptual distinction is well known to some Higher Order Theory of consciousness proponents. For example, Rose (2006, pp. 369–373) discusses the distinction between *Higher Order Perception* (HOP) and *Higher Order Thought* (HOT). HOP posits that something like a Global Workspace “scans” lower-order representations (“sensory impressions, phenomenal, creature consciousness,” p. 369) and perceives the content to which it attends. HOT (as per Rosenthal) posits that conscious content only arises when you think about unconscious lower-level representations. My three types of phenomenal content do not, however, depend on Higher Order Theories,¹⁸⁸ but merely on the more general and relatively uncontroversial idea of hierarchical cognition—the idea that some content is processed to produce more highly refined content—and that we can distinguish between conceptual and perceptual varieties of such higher-order content.

Another arena that utilises this conceptual/perceptual distinction is the debate over *cognitive phenomenology* (Bayne & Montague, 2011)—whether cognitive content must always be perceptual in character, or whether there is genuinely non-perceptual, purely abstract conceptual content. Again, I invoke this debate only as

¹⁸⁷ Zimmerman et al., (2016) draw a similar distinction in auditory experience between attention to higher-order features (Seeming-that) and attention to sensory information itself (Experience-of), and find distinct parietal pathways for each.

¹⁸⁸ In fact, the idea of phenomenal FOC may well be antithetic to some kinds of Higher Order Theories (9.5.2.1).

evidence for the validity of the distinction in theory at least, without wishing to enter into its finer points.¹⁸⁹

FOC and higher-order content do not necessarily imply a Duplicate Store or Partial Duplicate Store model (5.4.1). The same content cannot be *both* FOC *and* Experience-of at the same time.¹⁹⁰ Rather, some FOC is processed,¹⁹¹ altered, and refined by Executive Attention in WM so that it *becomes* Experience-of content. The original FOC no longer exists, much as half-way through a chess game, the original positions of the pieces no longer exist.¹⁹² There is evidence that the implementation of attentional strategies enhances processing of visual content, not only in the visual periphery, but even within the foveola itself (Poletti, Rucci, & Carrasco, 2017): i.e., FOC becoming Experience-of. Seeming-that can be thought of as a kind of pointer under a File Directory model. It is the propositional “tag” that describes other content, whether that content lies within or outside WM.

7.2.1.4 *Relevance to Overflow and Q*

My three types of phenomenal content provide the conceptual tools for characterising and perhaps identifying whether there is in fact phenomenal overflow. Experience-of and Seeming-that are the products of Executive Attention. If

¹⁸⁹ The idea of phenomenal Seeming-that would be rejected by those who reject cognitive phenomenology.

¹⁹⁰ See also the debate over whether WM recruits visual sensory processing areas to keep visual content in WM (Gayet, Paffen, & der Stigchel, 2018; Scimeca, Kiyonaga, & D’Esposito, 2018; Xu, 2018).

¹⁹¹ On attention modulating both conscious and subconscious content, see Sumner et al., (2006).

¹⁹² Incidentally, this provides a plausible account of Neander’s (1998) triplet whose Seeming-that “I have a sensation of green” follows from an actual first-order experience of seeing red. On my account, this is merely a case of faulty (non-veridical) higher-order processing. Even if the triplet has a phenomenal FOC experience of seeing red, it is quickly overwritten by an Experience-of—and consequently, a Seeming-that—she sees green, before she has time to form a Seeming-that she saw red. Compare the accounts of Rosenthal (2009, p. 249) and Block (2011b, pp. 423–424).

there is to be cognitive content that is phenomenal but not Executively Attended, it is in FOC that it will be found. The two particularly pertinent questions are:

[1]. Can FOC ever be phenomenal, independent of any higher-order processing?

[2]. How can we come to an answer to [1]?

The particular kinds of higher-order Executive Attentional content I have called Experience-of and Seeming-that are especially relevant to the epistemological question [2], since they are the standard means by which [1] has hitherto been answered. They constitute *report*, to which I return shortly.

We can now reframe the competing views of Block and Cohen et al., using this terminology as follows. Block can be taken to hold that the richness of phenomenal FOC is not limited by the capacity limitations of the higher order Experience-of and Seeming-that. Cohen et al., can be taken to hold one of two views: either that FOC is phenomenal, but is as sparse as Experience-of and Seeming-that; or that FOC is not phenomenal at all—only ensemble statistics that are averages of rich unconscious FOC become phenomenal as content in the higher-order Experience-of.

7.2.2 Access Consciousness and Executive Attention

There is a potential for some terminological confusion here that is worth clarifying from the outset. I take Block's (1995, p. 231) "access consciousness" to be roughly equivalent to my Executive Attention. He defines it with three sufficient but not necessary properties: content that is poised for use in reasoning, rational control of

action, and speech.¹⁹³ This overlays my concept of Executive Attention as structuring content for higher-order processing, although I have focused on perception (Experience-of) as well as conception (Seeming-that). This is in keeping with how others have understood access consciousness (Gross, 2018, p. 2; Lamme, 2003, p. 14). Keep in mind that access consciousness can be “non-phenomenal” (1995, p. 231)—indeed, it is one of Block’s chief interests to show that phenomenal consciousness can come apart from access consciousness.

7.2.3 Broad Definition of Report

The standard way of determining precisely what content is phenomenal for a subject is *report*, by which I mean more than just verbally reporting what one experienced. I use report in a broader sense, to encompass any higher-order phenomenal content by which a subject is aware of the phenomenal content of their experience. Thus, report includes Experience-of and Seeming-that, although it may not be exhausted by them (7.2.1.3).¹⁹⁴ Report may include verbal communication, but also non-verbal communication (e.g., pressing a button), and conscious knowledge that is poised for external communication, but is never actually externally communicated.¹⁹⁵ It seems obvious that we very often report (in this sense of “report”) the content of our

¹⁹³ Smithies (2011) presents a very similar account of attention as rational-access consciousness, although he seems less amenable to the possibility of dissociating phenomenal consciousness from access consciousness.

¹⁹⁴ Compare Sergent and Rees (2007) who argue that “conscious access overflows report” (p. 524). On my terminology, they are arguing for the view that phenomenal Experience-of overflows phenomenal Seeming-that (which my account allows), and against Block’s view that phenomenal FOC overflows both Experience-of and Seeming-that (which my account refutes, agreeing with Block).

¹⁹⁵ A definition roughly similar to that of Naccache (2018, pp. 3–4).

conscious experience. But it does not follow from that that all that is consciously experienced is reported, or even reportable.¹⁹⁶

As I use it here, report is always phenomenal. Thus, the involuntary dilation of one's pupils of which one is oblivious in response to an emotional stimulus does not count as report, but having the unverballed conscious realisation that I am scared does. As I use it, all report is phenomenal, but not all phenomenality is report. Inherent in one of my four characteristics of consciousness—what it is like-ness—is an idea very similar to report. The descriptive statement “what it is like” is very close to the question “what is it like?” Questions demand an answer, and an answer is very similar to a report. However, it may be like something for me to see a rose without my ever answering the question, “what is it like?”¹⁹⁷ Thus, report as I use it here is necessarily phenomenal, but not all phenomenality is necessarily report (or reported). There may be subconscious Experiences-of and Seemings-that, but they would not be report as I use it here.

Indeed, overflow dictates that there is content that is phenomenal—some part of FOC—but not reported,¹⁹⁸ and is not itself report. Report requires Executive Attention, since it is a kind of further processing by the cognitive executive that implements attentional strategies. If we are to find instances of overflow, it will be by finding FOC that is not accessed in these ways (Fazekas & Overgaard, 2018).

¹⁹⁶ For example, some vegetative patients exhibit normal fMRI patterns of activity when asked to perform mental imagery tasks. This would suggest that they are consciously experiencing the appropriate content for performing those tasks but utterly incapable of externalising that content in any way (Klein, 2017). Conversely, strictly speaking, externalisation of content is not itself a guarantee of consciousness—a record player playing a voice saying “I am conscious” is not (thereby) conscious. It is far from established that even very complex artificial intelligence delivers consciousness, even if programmed to answer questions about its consciousness in the affirmative. And in humans, there are experimental paradigms where implicit content makes its way into reporting.

¹⁹⁷ Compare Stoljar (2016, pp. 1182–1183).

¹⁹⁸ It never gets further processed into Experience-of content, although it may (or may not) Seem-that it is phenomenal to the subject. I elaborate on this below (8.6.2.2).

7.3 Florence and Ben

In this section I draw one more important distinction that has—as far as I can tell—been almost entirely overlooked in the literature on Q or overflow—the possibility of *foveal* rather than *peripheral* phenomenal overflow. Consider Florence, who is enjoying the view of a beautiful forest from her balcony. She stands perfectly still as she takes in the whole magnificent vista, savouring its grand complexity, texture, colour, and the gentle movement of the branches and leaves in the breeze. To Florence, it seems that she phenomenally experiences seeing every tree, perhaps even every leaf visible before her in the forest.

The overflow proponent argues that even if Florence does not have a rich Experience-of every tree, at least her phenomenal FOC contains every tree, and it is this that justifies her Seeming-that she has a visual experience of every single tree, even if she can't report the fine detail of every single tree. The overflow sceptic counters that Florence is mistaken in this Seeming-that—it is only an *illusion of richness*, arising from sparse ensemble statistics of the attentional periphery and a kind of perceptual assumption that the periphery must be as rich as the focus of attention.

There is one very good reason to side with the overflow sceptic in Florence's case. When Florence holds her head still and gazes upon the forest vista, the light from that scene strikes her retina. But her fovea is far richer in receptors than the rest of the retina. It also has various other features—pushing aside of ganglion and vascular cells, thinness, etc.—that maximise its resolution (Galvin, O'Shea, Squire, & Govan, 1997, p. 2035; Hall, 2015, p. 647; Schwartz, 2017, Chapters 2, 3). But away from the fovea, in the rest of the retina, the receptors are sparser, and the conditions less conducive to high-resolution vision. What is more, although the fovea represents 0.01% of the total area of the retina, approximately 8% of the striate visual cortex (primary visual cortex, or V1) is devoted to processing the information that derives from it. There is even evidence that it is more difficult to attend to the retinal

periphery, and that a number of factors contribute to poorer performance on various visual tasks when the stimuli are in the periphery rather than foveated (Staugaard, Petersen, & Vangkilde, 2016). If Florence is mistaken about it Seeming-that she richly experiences every tree, there may be a very good reason for that: that richness of detail does not proceed beyond the peripheral retina.

Ensemble statistics do an admirable job of explaining why Florence nonetheless feels that she does experience every leaf, richly, and why she never notices small changes in the periphery that do not significantly alter the value of those statistical averages (Greenwood, Bex, & Dakin, 2010). The brain fills in or interpolates content to make it Seem-that she experiences every tree and leaf, much as it fills in the receptor-free blindspot where the optic disc is located on the retina (Ramachandran & Hirstein, 1997, pp. 434–437).

Other strategies also contribute to this feeling of overall richness. The eyes generally do not stand still when we gaze upon a scene, but saccade rapidly and constantly without us even realising this is happening (Henderson, 2003; O'Regan & Noë, 2001), flitting to different parts of the scene so that the brain can take all of this rich information and weave it seamlessly into a broader tapestry. We also tend to *inflate* or overestimate the richness of peripheral visual content (Odegaard, Chang, Lau, & Cheung, 2018). All these strategies combine to explain why Florence might be mistaken about it Seeming-that she phenomenally experiences every tree richly. In fact—and the reader can try this at home—if she holds her eyes still and focuses her spatial attention on a region in her peripheral retinal field, no amount of attention will bring forth an experience of rich detail.¹⁹⁹ Information about average colour and brightness may be preserved and phenomenally accessible to her Experience-of (Saiki & Holcombe, 2012), but what is unavoidably lacking is rich, high-resolution visual phenomenal content. Cohen et al., explicitly state that their ensemble statistics model applies to extrafoveal visual perception:

¹⁹⁹ Compare the phenomenon of *crowding* (Block, 2013b, p. 131).

“We argue that items that are attended to and foveated are perceived at a higher resolution, while items that unattended or are in the periphery are primarily perceived as being part of an ensemble” (Cohen et al., 2016b, p. 327).

So, the same considerations do not apply to foveal vision.²⁰⁰ Consider Florence’s brother, Ben, who shares Florence’s love of nature. But whereas Florence’s tastes are heavily inclined towards flora, Ben’s incline more towards fauna. He is delightedly attending closely to a relatively small yet exquisitely clear photo of a speckled hen²⁰¹ on a page in front of him. When asked if he can see every individual speckle, Ben answers quite emphatically, yes. The hen fits completely within his foveal visual field.²⁰² Here, there is no retinal filtering of data in the sense described above, for the receptor density across the fovea—and therefore, across the whole speckled hen—is much more than adequate to resolve every speckle (Appendix 7). What is more, the whole of the hen fits easily within his focus of spatial attention—none of it is spatially peripheral to that focus—so long as he focuses his attention on the hen as a whole, rather than on any smaller part of the hen. And this is quite an easy thing to do, given that the whole hen fits within his foveal field. All that rich data—the detail of every single speckle—passes through the foveal retina to the early visual processing areas of the brain. Neither do his eyes need to saccade to take in more detail from a broad scene. By keeping his eyes perfectly still, Ben can foveate the whole hen. If Ben cannot truly experience each speckle, if his phenomenal FOC of the hen is sparse, it will *not* be because the content is absent from his visual processing centres, as is the case with Florence’s peripheral trees.

²⁰⁰ Although I highlight the contrast between Florence and Ben cases in this chapter and focus chiefly on establishing foveal overflow, this kind of analysis may profitably be applied in support of overflow in Florence cases too (see 9.4.2 for a brief sketch).

²⁰¹ See Tye (2009) for an account of the origins and history of the speckled hen in philosophical discourse, and some recent debates.

²⁰² For a technical analysis of the Ben’s visual anatomy and physiology, see Appendix 7.

The kinds of sceptical arguments raised by Cohen et al., will not avail in Ben's case, or at least, they require significantly more justification. Recall that Florence's peripheral content is sparser than her foveal content. Initially, it Seems-that she has rich peripheral phenomenal FOC, but when she is asked about it, she discovers that it is much sparser than she naïvely thought. When Florence "zooms out" her spatial attention to take in the whole vista or holds her eyes still and shifts her spatial attention to the periphery, she comes to doubt her initial Seeming-that she saw the whole forest and every tree and leaf in it clearly. She cannot in fact, by Executively Attending, generate a clear and distinct Experience-of all those individual items. The best she can hope for is filled in content based on ensemble statistics, which cannot deliver truly rich and veridical content, any more than if she were to close her eyes and imagine the scene.

But things are different for Ben. Like Florence, it initially Seems-that he has rich phenomenal FOC of every speckle. But when he brings Executive Attention to bear to answer the question of whether he sees every speckle richly, *he does not come to doubt it* as Florence does. Unlike Florence, he only becomes *more* confident in his Seeming-that he sees every single speckle clearly. However, like Florence, he is very limited in what he can *report* about those speckles. Despite seeing them clearly and richly (FOC) he cannot report on the number, individual shape or precise spatial relations (Holcombe, Linares, & Vaziri-Pashkam, 2011) of more than a few speckles at a time—just those speckles to which he individually turns his object Executive Attention. He cannot even subitise a small group of speckles, or even judge the left half of the hen to have more speckles than the right, without employing his Executive Attention. So Ben's Experience-of the speckles—the higher-order perception that is grist for those kinds of judgements—is sparser than his FOC.

What is more, if Ben maintains the hen in his foveal field but shifts his spatial attention to an object in the periphery, so long as his lenses maintain optical focus, he continues to see every speckle richly. In this case, there is a diminution in his Experience-of the speckles—the speckles play much less of a role in his voluntary stream of thought, become less *prominent*—but, I propose, no diminution of

phenomenal FOC. Neither the actual content nor the phenomenality as such are diminished.

Ben's incorrigibly rich and clear experience of every speckle, coupled with the relative sparsity of his reportable Experience-of the speckles seems like a clear case of overflow. The overflow sceptic has two broad ways of objecting. First, she might place the burden of proof on the overflow proponent to establish that phenomenal experience is capable of such high-capacity rich content in the first place—the topic of the next chapter. Second, she might argue that Ben (or we) have no way of *knowing* that Ben's foveal phenomenal FOC is rich—the topic of the balance of this chapter.

7.4 Epistemic Woes

When we come to identify whether or not phenomenal experience is rich or sparse, we are faced with a major epistemic hurdle.

“Reporting on what we see requires us to attend to it. So the failure to report an object of visual consciousness might reflect a failure to attend to the object, rather than an absence of visual consciousness of the object” (Stazicker, 2011a, p. 163).²⁰³

²⁰³ See also Block (2007, p. 483), Shea and Bayne (2010, pp. 478–479), and Philips (2018, p. 1).

The phenomenal overflow content in question is, on my distinction above, just a subset of FOC that is phenomenal, but not interpreted or further processed by capacity-limited Executive Attention in any way. In particular, it never gives rise to Experience-of or Seeming-that content, phenomenal or not. And here's the rub: our chief—perhaps currently sole—way of probing FOC is via report (Experience-of and Seeming-that). The only way to probe content that undergoes no further cognitive processing is via further cognitive processing. The only way to identify conscious content that is not Executively Attended is via Executive Attention. The only way to determine whether iconic memory content is phenomenally rich is via the apparatus of WM, which is phenomenally sparse. So, the issue turns upon epistemic justification: how do we (we, the investigators, or we, the subjects like Ben and Florence) confirm whether phenomenal FOC is rich or sparse, without relying on such higher-order content?

The overflow sceptic argues that phenomenal FOC is sparse because when I turn my Executive Attention to it, I find it to be sparse. But Executive Attention is capacity limited, due to the nature of WM (Chapter 8). Yes, the report produced by Executive Attention is itself sparse, but that would be the case whether the FOC it purports to report is sparse or rich, simply because Experience-of and Seeming-that—the bearers of the report—are by their nature sparse. This is akin to taking a photograph of a bold and vivid Picasso or a Warhol with an old, grainy, low-resolution camera and holding the original painting to be grainy and smudged on the evidence of the photograph. If the only evidence we have available is the grainy photograph, it is as plausible that the original painting is clear and vivid as it is that it is grainy. In fact, if we *know* that the camera produces grainy images even when the original scenes are known to be clear and vivid, then we have additional reason to be cautious about coming to the conclusion that the original painting was grainy. That is the case here. We *know* that WM is capacity limited, so we should be cautious about imputing the same limitations to FOC which, on my account, lies outside the capacity limited Executive Attention of WM.

In this section, I develop this argument by first invoking the diachronic nature of cognition to show that over time, we can only report a tiny subset of that which we

in fact experienced. This principle can be telescoped down to as small a time-period as we like and still hold. I then argue that the intuitive immediacy of phenomenal experience provides powerful warrant for taking the intuitive richness of phenomenal FOC more seriously. Finally, I consider the promise of no-report paradigms for getting around the epistemic impasse of reliance on report, but find it to be fatally flawed.

7.4.1 Synchronic or Diachronic Experience?

In 2.4.4 I counted temporality as a core feature of phenomenal consciousness. We commonly take our synchronic conscious experience to be something like a snapshot in time, a discrete frame on a film, completely dissectible from all other moments in the stream of conscious experience. But the reality is that conscious experience—and therefore phenomenal content—is fluid. Consider the case of reading a novel. It takes many hours to read the novel, perhaps spread out over days or weeks. Each sentence I read in that novel is rich in content, striking connections in my mind to memories, perceptual imagery, emotions, and imaginations. These are not only experienced at the time, but also attended. Yet there is also a vast body of experiential content, what Dainton (2000, p. 31) calls the “phenomenal background”—stray background noises, rumblings of one’s tummy, gasps at plot twists—that is only slightly or not all attended as I read the novel. By the time I finish reading the novel, all I have left is an extremely sparse Experience-of the detail of the story, such that I will only be able to report the smallest subset of all that I experienced—events, characters, my reactions—over the time I was reading the novel. Coupled with this, however, is an ostensibly reliable Seeming-that I did

indeed experience all the richness of the novel and perhaps the background over time, despite the sparsity of my report about its details.²⁰⁴

Now something very similar might well be happening to all of us on a much shorter timescale all the time. Ben is looking at the hen. Specific aspects of the FOC are being constantly interpreted by Executive Attention into Experience-of and Seeming-that. All three are constantly being bound into a unified phenomenal experience. When Ben comes to report the content of his experience, however, he cannot simply report his phenomenal FOC precisely and completely at this present moment. Executive Attention takes time, however brief, to do its work. Like the novel, Ben is always reporting interpretations of *past* FOC. And like the novel, much of the rich content in iconic memory is constantly frittering away, with only a small proportion of it being preserved from moment to moment by rehearsal and manipulation in WM (8.2.4) and making it through to conscious report. Even that preserved content is constantly competing with new preserved content, such that any particular preserved content is preserved in consciousness for a limited period of time, before it is replaced by other preserved content (Baddeley, 2007, Chapter 7; Oberauer, Farrell, Jarrold, & Lewandowsky, 2016). Even synchronic experience is dynamic, and therefore in a sense, diachronic.

This inherent diachronicity, together with the well-established fact of the fragility of rich very short-term stores (such as iconic memory) ought to give us serious pause with respect to overflow-sceptical arguments based on the illusion of richness. A synchronic visual experience is really a temporal process, and by the time Experience-of or Seeming-that is produced by this temporal process, much of the FOC has already frittered away beyond retrieval for any kind of further processing. The nature of capacity limitations of the mechanisms of report (Chapter 8), taken together with the fragility of very short-term stores, suggest that whether the FOC is phenomenally rich or not, we would face the same sparsity of report. The sparsity of

²⁰⁴ Ward (2018, pp. 3–4) makes a similar point on a much shorter timescale—the RSVP paradigm, where a series of letters is flashed quickly before a subject, who experiences every letter, but cannot report most of them.

report, then, cannot be taken to be a reliable guide as to the sparsity of phenomenal FOC.²⁰⁵

We need some other way to decide whether the phenomenal FOC is sparse. In the case of a someone reading a novel, we confidently trust her Seeming-that her diachronic experience of the novel was rich, despite her current sparse Experience-of the novel. The burden of proof falls upon the sceptic to give good reasons to doubt that Seeming-that. Perhaps, then, by analogy, the burden of proof falls upon the overflow sceptic to show why we should not, in the appropriately parallel case, trust Ben's Seeming-that he phenomenally experiences every speckle despite his sparse Experience-of the hen?

In the next section, I invoke the immediacy of phenomenal FOC to argue for the reliability of Ben's Seeming-that he sees the speckles richly, which has the effect of strengthening the overflow proponent's position. In the section following that, I consider another way to address this epistemic obstacle: bypassing report altogether via no-report paradigms, but find in these a fatal flaw. This has the effect of diminishing our hopes that this question will ever be finally settled by this route.

7.4.2 Immediacy

“Without further argumentation, phenomenology cannot serve as the sole basis for any theory of reality. It may be taken only as a finger, *pointing* in some direction, rather than *conclusive evidence*

²⁰⁵ Space prevents me from delving deeper into the implications of the tantalising finding that cognitive content can be stored for a few seconds *without neural activity* (Trübutschek et al., 2017), so-called “activity-silent WM,” and whether this might instead represent Executive-Attention-free phenomenal perceptual content.

for or against a particular thesis ... But we must be careful here, for first-person reports can also be veridical or even sources of wisdom" (Forman, 1999, pp. 619, 620, italics in the original).

In this section I argue that there is a very powerful reason to take foveal phenomenal overflow seriously: the *immediacy* of experience.²⁰⁶ To Ben, it clearly and confidently Seems-that he sees every speckle richly. He is at least as confident about this as he is about his rich experience of the novel he read last month. And Ben's increased confidence sharply contrasts with Florence's loss of confidence that she sees the whole forest richly, when they both introspect their own level of richness more carefully. This datum, while by no means conclusive, must be taken seriously.

At the heart of the overflow proponent's view, I suggest, lies an intuition that Ben phenomenally experiences every speckle *immediately*, in the most literal sense of the word: without the mediation of reflection or Executive Attention.²⁰⁷ This is the case whether or not he further processes that content to produce Experience-of or Seeming-that he saw the speckles. Carruthers captures this "powerful intuition:"

"our awareness of our own experiences is *immediate*, in such a way that to *believe* that you are undergoing an experience of a certain sort *is* to be undergoing an experience of that sort" (Carruthers, 2016, sec. 3, italics in the original).²⁰⁸

²⁰⁶ Compare Kriegel's (2007, p. 132) *Method of Knowability*, and his idea of "first-person knowledge," characterised as effortless, without the mediation of inference, and quick, although he is thinking more of the contrast with "third-person knowledge."

²⁰⁷ To be quite clear, I am not thinking here of the immediacy of perception of *external objects*, as per Locke, Berkeley, and Hume (Shoemaker, 1994, p. 251).

²⁰⁸ Similar ideas may also be found in the fourth century Augustine (1950, sec. 3.11.24) and in Descartes (1637).

Ben can be mistaken in the veridicality of what he sees—his experience of a hen might be purely a hallucination. Ben might be mistaken when he reflects upon what he sees—he might miscount the number of speckles or be fooled by a clever illusion that makes the hen look like a Heffalump. But these mistakes occur in the subconscious processing prior to FOC or the higher-order processing posterior to FOC. The phenomenally experienced FOC is itself *immediate*, and therefore not liable to error, *in itself, qua experience*. As Horgan (2012, p. 406) puts it, “there is no gap between appearance and reality, because the appearance just is the reality.”

Metzinger (2003, p. 125) discusses a similar concept, *phenomenal immediacy* in relation to representationalism. But my use here of the term *immediacy* differs from *scene-immediacy* (Schroer, 2012) and has little to do with representationalism or naïve realism. It is not about the subject’s relation to the object of experience out there, in the world (Lyons, 2017; Masrour, 2011, pp. 381–383). It is about the reflexive relation between the subject and the phenomenal FOC of experience. When Ben looks at his hen, not only is the “I-ness” intrinsic to the experience, so also is the FOC.

There is something it is like for Ben to see the hen as a whole—immediately, without reflection, analysis, contemplation, report, or any other kind of higher-order processing of that content. But there is also, for the overflow proponent, something it is like for Ben to see every speckle on the hen—immediately, without reflection, analysis, contemplation, report, or any other kind of higher-order processing of that content—and hence, for many of those speckles, without Executive Attention being involved.

The overflow sceptic must reject this latter possibility.²⁰⁹ Ben must be mistaken precisely about the richness of his phenomenal FOC. If his phenomenal FOC of every

²⁰⁹ I explore the different ways this rejection can be framed in 8.6.

speckle is richer than his Experience-of the hen, then phenomenal content overflows attended content. The capacity limitations of Executive Attention and WM guarantee that his Experience-of cannot be so rich. It may—by virtue of ensemble statistics in the periphery of attention and the nuances of WM capacity (8.2)—be in some sense richer than the traditional three-four items of WM capacity, but it cannot be as rich as the overflow proponent would have it (Cohen et al., 2016b).

There seems to be a contradiction here, between the apparent sparsity of Experience-of on the one hand, and the Seeming-that Ben's experience is rich, based on the immediacy of Ben's phenomenal FOC experience, on the other. Is there any way of determining which of the two kinds of report is the more reliable for our purposes? I consider first some positive reasons for trusting Ben's immediacy-motivated Seeming-that his phenomenal FOC is rich, before considering some reasons for doubting it.

7.4.2.1 *For the Veridicality of Immediate Richness*

"The intrinsic phenomenal aspects of one's current experience are epistemically special because they are self-presenting and are not subject to any appearance/reality gap" (Horgan, 2012, p. 420).

Odegaard et al., (2018, p. 8) observe that while a subject's "confidence is not synonymous with phenomenology *per se* ... there are many cases where confidence provides an effective assessment of phenomenology's presence or absence." They give the example of blindsight, where the absence of the subject's phenomenology is taken as veridical despite the preservation of the subject's ability to perform tasks using the unconscious content. We might think about Florence and Ben's confidence in terms of the concepts of "feeling of knowing" (Hanczakowski, Zawadzka, Collie, & Macken, 2017; Mangan, 2001) or "feeling of rightness" (FOR) (V. A. Thompson,

Prowse Turner, & Pennycook, 2011) and the “feeling of error” (Gangemi, Bourgeois-Gironde, & Mancini, 2014). While these feelings are by no means infallible, there are certainly situations where they are eminently trustworthy, as in the case of feeling that you know the sum of two plus two to be four. We evaluate our feelings of confidence on the basis of additional evidence, but Ben’s case is one where—due to its very nature as first-person experience—it is very difficult indeed to garner *third-person* evidence (7.4.3). Nonetheless, there follow two arguments for the veridicality of Ben’s immediate and confident Seeming-that he sees every speckle richly: Ben’s increasing confidence on probing; and the veridicality of his Seeming-that in tracking the dynamics of iconic memory.

7.4.2.1.1 *Confidence*

First, as observed above, Ben’s confidence in his Seeming-that he sees *every* speckle richly only increases when he introspects his experience more closely, in a way that Florence’s does not.²¹⁰ There is something *irrevocable* (Ramachandran & Hirstein, 1997, pp. 437–438) about Ben’s experience—his Seeming-that he sees every speckle richly cannot voluntarily change, any more than his Seeming-that two plus two is four. Even so, when Ben Executively Attends to the hen, he cannot, no matter how he tries, generate an Experience-of *every* speckle richly. He cannot hold in WM and manipulate the exact shape, number, and spatial relations of every single speckle in the way that he could, say, just three small dots. What Ben *can* do is Executively Attend to the hen and generate a confident Seeming-that he richly experiences every speckle (as FOC). Importantly, this Seeming-that is *propositional*—it “describes” the richness of his experience, but is not itself rich, although it is irrevocable. For Ben, there are no lacunae, and he can report on any single speckle (or few speckles)

²¹⁰ Ramsøy and Overgaard (2004, p. 10) argue that “one cannot *a priori* assume that reports of certainty work as reports of awareness.” Below I argue that there are good *a posteriori* reasons for assuming that Ben’s reaffirming certainty does work as a report of his awareness, especially by contrast to Florence’s diminishing certainty.

richly whenever he wants to, by further processing FOC content into Experience-of content.

Florence can't do this. When she Executively Attends to her forest—when she, holding her eyes still, zooms out her spatial attention as widely as possible, and/or shifts her spatial attention away from the fovea, her Seeming-that she sees every tree and leaf richly is quite revocable—she becomes more confident of the *sparseness* of the phenomenal content of her peripheral visual field for the reasons outlined above. Ben's individual speckles are *immediately* phenomenally experienced in a way that Florence's individual trees never can be (unless they are foveated).

Is subjective confidence a good measure of the consciousness of content? Cosmelli et al., (2007, pp. 761–762) propose the “structural invariants approach” to improve the reliability of first-person report. Basically, details of report that are consistent across different individual subjects are more reliable. I can find no formal empirical testing of the replicability of Florence and Ben's evolving confidence levels (it would make a worthwhile research project) but my own anecdotal research confirms their replicability. Pinto et al., (2017, p. 217) describe a principle whereby conscious content can be identified:

“Generally, it is thought that if subjects provide correct answers but feel that they are guessing (low confidence), their performance may be based on unconscious information. However, if subjects are confident when they give correct answers and feel that they are guessing only when they give inaccurate responses, then the information underlying the performance is thought to be consciously available (Merikle, 1992; Pasquali, Timmermans, & Cleeremans, 2010).”

For Florence, her initial confidence in Seeming-that she sees every tree and leaf is shaken by further introspection so that she now comes to think she might have been guessing after all, suggesting that she may not have been conscious of that rich detail. Phenomena such as change blindness, inattention blindness, and crowding are all examples of this principle at work, albeit in different ways or for different reasons. They all illustrate the fallibility of our reporting outside the focus of attention, but crucially are accompanied with that feeling of guessing, if not initially, certainly on further introspection or on being confronted with the reality (e.g., after being shown the gorilla video without counting basketball passes). Expectation can create illusory content (Aru & Bachmann, 2017; Mack, Clarke, & Erol, 2018). But Ben's introspection not only leads him to grow more confident in his Seeming-that he sees every speckle richly, his confidence is confirmed by the objectively richly speckled hen. While there remain some questions about the validity of these kinds of measures, they do represent positive evidence that Ben is indeed conscious of every speckle richly in a way Florence is not.

Empirically, little if any work on this question has been done on Ben cases. Naccache (2018) relies on explicitly parafoveal experiments to argue that taking introspective reports as veridical is gullible (p. 3). Odegaard et al., (2018; 2016) also address Florence cases, arguing that misplaced confidence in what one sees increases with distance from the fovea (Odegaard et al., 2018, p. 8). What is more, the kind of misplaced confidence exhibited by subjects in that study is exactly the kind that is defeasible by further attention of the right sort, as when Florence tries to ascertain whether she really can see individual trees and leaves in her visual periphery, and realised that she was mistaken in her initial Seeming-that.

7.4.2.1.2 *Tracking Iconic Memory*

Second, Ben's Seeming-that reflects the way iconic memory works in a way that suggests that it is a reliable indicator of his phenomenal experience. Haun et al., (2017, pp. 2–3) argue convincingly that the simple discriminations tasks or forced choice reports that dominate the empirical work on these matters fail to capture the

richness of daily perception, and suggest more creative ways of identifying that richness (such as asking an artist to draw as many details as they can remember of an image). If Ben gazes even briefly at his hen, then closes his eyes and immediately attends intently to the persistent traces of this visual image, he will find that for some fraction of a second he continues to be confident in it Seeming-that he sees the *afterimage* of every speckle richly. But this only lasts a fraction of a second, after which he loses even the Seeming-that he has maintained a rich afterimage of the hen. The afterimage quickly degrades and becomes sparser before fading altogether.²¹¹ I encourage the reader to try this at home.

Here, Ben's Seeming-that is accurately tracking the dynamics of iconic memory (and longer-term memory as well) and revealing to him the *difference* between the richness of the direct perception and initial afterimage compared to the rapidly degrading richness as the afterimage fades from iconic memory. This accurate reflection of the behaviour of iconic memory is important third-person evidence that Ben's confidence in his Seeming-that he sees every speckle richly is well founded.

7.4.2.2 *Against the Veridicality of Immediate Richness*

The overflow sceptic must argue that Ben's Seeming-that he experiences every speckle richly without Executive Attention is mistaken. Ben's Seeming-that is a form of introspection, and there are much discussed problems with introspection (Schwitzgebel, 2011; Smithies & Stoljar, 2012). There are a number of objections the overflow sceptic can press here: that the evidence for Ben's Seeming-that is too scant; that Seeming-that is generally unreliable; and that no speckle is in fact experienced without Executive Attention.

²¹¹ Perhaps a very ghostly and sparse image can be maintained for minutes, or even encoded into long-term memory.

7.4.2.2.1 *Lack of Evidence*

First, the overflow sceptic might argue that the admission in the last section that it is very difficult to find direct independent evidence of rich phenomenal FOC means there is a lack of direct evidence for immediacy arguments, which weakens them considerably. In response, I repeat that I am pursuing an abductive argument here, and the evidence presented in the previous section, while not voluminous, is: (a) powerful indirect evidence for foveal overflow; and (b) the best one could hope for given the impenetrably private first-person nature of overflowing phenomenal FOC (see 7.4.3). Whether overflow is real or not, this is the best evidence we could currently hope for (always leaving room for future novel approaches).²¹² The immediacy of phenomenal FOC in particular demands an alternative account from the overflow sceptic.

7.4.2.2.2 *Seeming-that is Fallible*

Second, the sceptic may argue that Seeming-that can and does get things wrong. In general, propositions—even about our own phenomenal content—require epistemic justification, a notoriously difficult area (Silins, 2015).²¹³ As I observed above, the reliability of introspection has been seriously questioned in recent times. Ben’s confident Seeming-that he sees every speckle richly could be a quirk of the brain, producing a powerful but unfounded *Feeling of Knowing* (7.4.2.1). This might arise via an interplay between sparse content and one’s expectations, filling in the gaps to produce the illusion of richness (de Gardelle, Sackur, & Kouider, 2009; Kouider, de Gardelle, Sackur, & Dupoux, 2010, p. 306), or merely an inflated metacognitive sense of richness (Odegaard et al., 2018). There are certainly other cases where Seemings-that fail to be veridical. Florence’s pre-reflective confidence

²¹² I make some further points on this topic in 8.6.1 and 9.5.2.3.

²¹³ For some fascinating examples of how wrong Seeming-that can be in pathological cognition, see Bayne (2010, pp. 218–219).

in the richness of her peripheral vision is one. Hypnosis is capable of making it Seem-that one remembers something they have actually forgotten (“hypnotic hypermnesia”) or forgotten something they actually do remember (Kihlstrom, 2007, p. 459). False memories can often be compelling (Brainerd & Reyna, 2005). Confidence and accuracy on perceptual tasks have been shown to dissociate (Koizumi, Maniscalco, & Lau, 2015, p. 1305).

But in all these cases, there are ways for the subject or the investigator to probe the Seeming-that and invalidate it. That is not the case for Ben, whose actual ongoing phenomenal experience of the speckles is continuously probe-able (unlike memories of experiences in the distant past), and for whom more probing leads only to more confidence that he sees them richly. Neither do we have access to direct²¹⁴ objective measures of the accuracy of Ben’s Seeming-that as we do in perceptual task paradigms.

On the other hand, there is considerable evidence that in some cases, Seeming-that *is* highly reliable, such as confidence ratings by patients with blindsight (Odegaard et al., 2018, p. 8). In fact, a tentative argument could be made that Ben’s rich experience falls under Horgan’s (2012, pp. 407–412) class of utterly “infallible introspections,” but that failing, it certainly meets all the criteria for his class of “super-reliable introspections” in which fallibility only occurs for things like a lapse in conceptual competence.²¹⁵ Taking all that into account, assuming the veridicality of Ben’s Seeming-that ought to be the default position until we have good reason to doubt it.²¹⁶

The sceptic can try to launch a Florence-style argument for the fallibility of Ben’s Seeming-that thus: the foveated speckles may not be in the periphery of the *visual*

²¹⁴ The tracking of iconic memory behaviour (7.4.2.1.2) is indirect objective evidence.

²¹⁵ For a succinct summary of the problem (*sans* solution) of where to draw the line between reliable and unreliable introspection in empirical paradigms, see Cosmelli et al., (2007, p. 763).

²¹⁶ Lamme (2010, p. 210) discusses what would constitute proper “confirmation” of phenomenality in such cases.

field, but they may be in the periphery of the field of *spatial attention*, if that is smaller than the foveal field. Thus, the same considerations would apply to the speckles as to Florence’s peripheral trees. But this will not work either. Ben can easily zoom his spatial attention out to encompass or even exceed his foveal field. And the specifications of the visual system that lent strong support to overflow scepticism in Florence’s case do not apply in Ben’s purely foveal case (7.3).

7.4.2.2.3 *It’s Not Really FOC (Executive Attention is Involved)*

Thirdly, the sceptic might concede the veridicality of Ben’s immediate rich phenomenal experience of every speckle, but argue that this is not actually FOC at all. This rich content only becomes phenomenal *by* being further processed in some way by Executive Attention such that it is actually phenomenal Experience-of. I consider this kind of response in detail in 8.6.2 where I explore the Expanded Attention account, on which the capacity of Executive Attention is large enough to encompass any amount of rich phenomenal content, thus negating overflow. For various reasons discussed there, this objection will not work either.²¹⁷

7.4.2.3 *Interim Conclusion on Immediacy*

There is, then, admittedly limited but compelling evidence for overflow from the immediacy of phenomenal FOC, and none of the sceptic’s objections seem to succeed. It has been proposed that “the claim that visual awareness overflows cognitive access must be supported by specific examples of visual input that can be consciously perceived without being attended, held in working memory, reported, or used to guide volitional action” (Cohen et al., 2016b, p. 332). I have argued that when we break down Ben’s experience to phenomenal FOC, Experience-of, and

²¹⁷ My discussion of higher-order theories of consciousness in 9.5.2.1 is also relevant.

Seeming-that, and when we take into account the inescapable datum of the immediacy of Ben's rich FOC experience of every speckle, the case of Ben's hen satisfies these criteria. The rich content itself is not held in WM, but it is reported on veridically—until we have good reason to think otherwise—as being rich.

But it would be nice to have a way to settle this question decisively. Is there any way for us (or Ben, for that matter) to circumvent the epistemic impasse posed by reliance on report and objectively determine whether his rich phenomenal experience of every speckle is veridical?

7.4.3 No-Report Paradigms and Neural Signatures

“But in fact we don't know what we are conscious of ... Once we acknowledge that, it is clear that we need other evidence about the presence or absence of a conscious sensation than introspection or behavior” (Lamme, 2010, p. 204).

What is required to circumvent the epistemic impasse described in this section is some way to objectively identify the contents of phenomenal FOC in detail without relying on report. This is the promise held out by *no-report paradigms*. Roughly, no-report paradigms are empirical paradigms that aim to identify cognitive content by means other than report. For example, it has been suggested that recurrent neural activity is the signature of consciousness (Lamme, 2010).²¹⁸ If such activity is measured in, say, the posterior visual centres of the brain, that is evidence that the visual content associated with that activity is conscious, even in the absence of

²¹⁸ The theory of recurrent neural activity is far from established (Tsuchiya & van Boxtel, 2010).

heightened activity in certain parieto-frontal areas (Lamme, 2006), generally thought to be correlated to Executive Attention. Simplistically, this suggests the *C~A* of overflow.

A more detailed example of a putative no-report paradigm is in Tsuchiya et al., (2015, p. 758): “In these cases it is important to ask to what extent we can be sure that such decoded contents reflect conscious experience – as opposed to unconscious stimulus processing.” They take Frässle et al., (2014) to illustrate a method by which conscious perceptual content can be identified without introspective report. Frässle et al., used physiological indicators, “optokinetic nystagmus and pupil size to objectively and continuously map perceptual alternations for binocular-rivalry stimuli” (p. 1738). Combining these with fMRI and introspective report enabled them to build a profile of what was happening objectively in the subject—physically and physiologically—when that subject’s perception of a bistable figure flipped from one image to the other. In theory, reading back from the objective signs ought to indicate which perception the subject is phenomenally experiencing. Another instance of a no-report paradigm cited by Tsuchiya et al., is Wilke et al., (2009), where once again, objective measurements of neural activity are correlated to the perceptual suppression of images in macaque monkeys.

More recently, Schelonka et al., (2017) suggest that distinct ERP signatures can be correlated with conscious awareness of content in an inattention blindness paradigm. Subjects went through three phases in this experiment. First, they performed a distractor task while words, consonant strings, or meaningless curves flashed briefly in the background. They were then probed for their awareness of the words, and as expected, 40-50% had failed to notice them. The second phase involved them performing the same kind of task after the awareness probing, so they now had some expectation that the words would appear in the background. Awareness probing after the second task showed—again, as expected—much improved recognition of the words. The third phase shifted their focus from the distractor task to the words themselves, and thus constituted a kind of control condition. Schelonka et al., measured ERPs for these three phases and claim that the

difference in these signatures between the first phase and the others is itself the signature of conscious awareness, and the difference between the third phase and the others is the itself the signature of report.

Pitts et al., (2018, p. 7) remark that “To date, very few no-report paradigms have been developed.”²¹⁹ The examples above illustrate some serious issues that are common in all of them, and likely to afflict any new ones that arise in future, which I now describe. I call them the *resolution problem* and the *correlate problem*. I conclude the section with some strong reasons for pessimism that any kind of no-report paradigm will ever be able to overcome these obstacles.

7.4.3.1 *The Resolution Problem*

First, existing paradigms paint with far too broad a brush to really settle the issue of foveal phenomenal overflow. Knowing whether a subject is experiencing one bistable image or the other (Frässle et al.) is a very different thing to knowing how much detail of either figure she is consciously experiencing.²²⁰ Similarly, recognising a word (Schelonka et al.) does not tell us much about the richness of phenomenal FOC. What would really help would be a method that picks out content more precisely. What we need is something that will tell us whether this specific speckle is being richly experienced by Ben.²²¹ While techniques are being developed to “read” cognitions from EEG or fMRI data (e.g., Nemrodov, Niemeier, Patel, & Nestor, 2018), we are still very far from identifying that level of detail in the brain.

²¹⁹ See also Sergent and Rees (2007, p. 524) for citations to earlier work in this direction, Naccache (2018, p. 4) for a discussion of some recent no-report experiments, and Storm et al., (2017) for a detailed survey and thoughts on future directions. For an example of no-report paradigms for identifying spatial attention, see Yeshurun (2019).

²²⁰ See a similar objection in Phillips (2018, p. 5).

²²¹ Compare the “sensitivity criterion” in priming (Doyen et al., 2014, pp. 15–16).

7.4.3.2 *The Correlate Problem*

Second, and much more seriously, even if we attained a high enough resolution, there would remain the problem of determining what exactly the signature correlates to.²²² The signatures in Frässle et al., distinguish between which of the rival perceptions is being experienced. But binocular perceptions supplant their rivals *via Executive Attention*. The signature may therefore be of Executive Attention, and therefore, of no use in searching for $C \sim A$.

Schelonka et al., promises to overcome this problem, since they infer discrete signatures of consciousness and report. But even here, a serious problem remains. It seems impossible to establish the correlation between a physiological indicator and specific conscious content without *in the first instance relying on report* to do so. And if so, we are just identifying the physiological marker of *reported* content, not unreported phenomenal FOC. If so, we cannot use such markers as indicators of *unreported* phenomenal FOC (Stazicker, 2011a, p. 164).

To identify phenomenal content without report, we need a neural correlate of consciousness (NCC) that is independent of report. But to determine what a NCC is in the first place, we must rely upon report. This paradox has recently been observed by a number of researchers regardless of their position on overflow (Bronfman, Brezis, Jacobson, & Usher, 2014, p. 1402; Kanai et al., 2006, p. 2335; Mogensen & Overgaard, 2018; Phillips, 2018; Pitts et al., 2018; Stazicker, 2018). Even the cleverest of paradigms are unable to overcome this obstacle (Pitts, Metzler, & Hillyard, 2014; Pitts, Padwal, Fennelly, Martínez, & Hillyard, 2014).

²²² There are also methodological issues. For example, de-Wit et al., (2016) point out that imaging studies often focus on how the *experimenter* interprets the data rather than how the subject's *cortex* interprets the data. These are two very different senses of "information" in the brain. Their discussions of the various pitfalls involved in trying to correlate the two is sobering.

Of course, for the purposes of addressing the overflow question, signatures need not themselves be the substrate of consciousness. It would be enough to identify any kind of signature—chemical (neurotransmitter), electrical (neuronal firing patterns), anatomical (a region of the brain, probably pretty small), or behavioural (a twitch of the left pinkie finger)—that *invariably* correlates to the phenomenality of the FOC content. But even so, the obstacle remains steadfast: to calibrate *any* kind of signature at all as invariably correlating to the phenomenal FOC, we must rely in the first instance on report. It seems that—unless there is some other way of determining NCCs without report—this objection is fatal to the hopes of identifying overflow via neural signatures.²²³

7.4.3.3 *The Importance of the Subjective*

My arguments for the immediacy of phenomenally rich FOC are abductive rather than deductive, and no-report paradigms seem incapable of circumventing our reliance on report. How might the quest to find $C \sim A$ in overflow proceed? One view I find unconvincing is that the question can only be settled in strictly objective terms:

“Currently, no experimental results uniquely support the existence of consciousness independent of function and access. Could future experiments accomplish this? We argue that all theories of consciousness that are not based on functions and access are not scientific theories” (Cohen & Dennett, 2011, p. 361).

²²³ Brief arguments for roughly this conclusion have been made elsewhere (Overgaard, 2010; Overgaard & Fazekas, 2016).

Later, they ask, “What does it mean to study consciousness without function?” (p. 362). In context, by function they mean accessing the phenomenal content for reporting, action, etc., i.e., higher-level processing. Cohen and Dennett’s preferred approach is to define consciousness *in terms of function*. The assumption here is that consciousness is known *only* in being so accessed. Thus, they also ask, “What does it mean to have a conscious experience that you yourself do not realize you are having?” (p. 362). By contrast, none of the four characteristics in my definition of consciousness (2.4) intrinsically involve function. They may certainly interact with function in some way, but they stand apart from it.

Perhaps this is the nub, or at least one crucial factor, that separates the overflow proponent from at least some kinds of sceptic. Cohen and Dennett’s stipulation that there is no consciousness without function seems to me to echo the mistakes of Behaviourism (3.3.2), imposing an unwarranted kind of verificationism where it is not appropriate. What is more, I argued in 3.3.3 that Phenomenal definitions of attention beg the question against Q, because they analytically define consciousness into attention. Cohen and Dennett are similarly begging the question by analytically defining function (Executive Attention) into consciousness rather than leaving it to empirical investigation.

Objectivity is a noble principle in scientific inquiry, and Cohen and Dennett’s call for an authentically “scientific” approach (p.361) is admirable, but this is one area where we need to broaden what constitutes “scientific.” Phenomenal consciousness is a strange beast indeed. Its first-personal character is quite possibly unique in the natural world. To expect the kind of third-person verification or falsifiability that is the standard scientific staple to handle it may be unrealistic. Were they to apply their stance to the intractable Problem of Other Minds (Hyslop, 2018), they would have to admit that they cannot be certain that anyone else is conscious. There is no escaping the simple, inescapable fact is that no one has direct access to my phenomenal content except me. Kriegel (2007, pp. 124–125) raises a patently false hypothesis: olfactory perception is not phenomenal. How do we know this hypothesis is false? Ultimately, it comes down not to any third-person evidence or

argument, but to first-person experience. The same might be said of how we know that proprioception *per se* has no phenomenal character (6.2.2).

How then might the quest to find $C \sim A$ in overflow proceed? In this chapter, I have argued on the grounds of immediacy for taking seriously Ben's Seeming—that his phenomenal FOC is rich. There is a kind of falsifiability here—it is just not *third-person* falsifiability. Whether that qualifies as “scientific” or not may not really be the issue, if one is willing to entertain that not all knowledge is scientific knowledge in the narrow, formal sense of “scientific” (which I am). Perhaps this is a question (like the Problem of Other Minds) than simply *cannot* be answered in the third-person, only in the first-person. A more promising direction, then, will be to improve the accuracy of and resolution of methods for assessing the veracity of subjective reports (Rausch, Müller, & Zehetleitner, 2015).

7.5 Chapter Summary

To summarise, our consideration of Florence and her forest found good neurophysiological and phenomenological reasons to think that phenomenal overflow is an illusion of richness in her case, as suggested by the overflow sceptic. However, those considerations find no traction in the case of Ben's hen. The immediacy of the richness of Ben's hen experience and the way it reflects the objective behaviour over time of human memory systems give us reason to favour overflow, and the objections raised thus far have serious flaws. Neither can objective methods answer this question, since they are too coarse, and suffer a seemingly fatal methodological flaw—they rely on report to be calibrated in the first place. We are left without a reasonable way to explain away Ben's confident Seeming—that he richly phenomenally experiences every speckle, even if he can't report on each speckle. Genuine $C \sim A$ seems to be the best explanation thus far.

But there is more evidence that can be brought to bear. In the next chapter, I underwrite the case for overflow from a consideration of the different capacity limitations on attention, consciousness, and WM.

8 **Triangulating Capacity Limitations**

8.1 Comparing Capacities

If consciousness is to overflow Executive Attention in WM, then it must have a relatively higher capacity for content. Overflow in the case of Ben's hen requires that the Target of consciousness—all the individual speckles with all their features—be phenomenal, while at the same Timing, the Target of Executive Attention be limited to only a small subset of those speckles and their features. The Capacity Question (5.3.2.3) has been studied extensively with regard to WM, and to a slightly lesser degree with regard to attention, but much less with regard to phenomenal consciousness.

In this chapter, I argue that the capacity of consciousness is significantly greater than the capacity of Executive Attention in WM. What is more, they differ in both nature and neural substrates. If this is so, I conclude, it would be remarkable indeed if their capacities were identical in a way that denies the *C~A* of phenomenal overflow. After a criminally brief overview of the current state of thinking about the

natures and magnitudes of the capacities of WM, Executive Attention, and phenomenal consciousness, I “triangulate” the three to develop an empirically informed schema—my *Witches’ Hat Model*—for how they interact. As mine is essentially an abductive argument—an inference to the best explanation—I conclude the discussion by considering some alternative explanations, none of which seem preferable to the one presented here.

8.2 The Capacity of Working Memory

That WM is capacity limited is a matter of everyday experience, but determining the nature and magnitude of this capacity turns out to be surprisingly difficult. Miller’s seminal paper (1956) argued on empirical grounds that there is a relatively rigid figure of 7 ± 2 items²²⁴ that can be stored and manipulated in what we would today call WM. And Baddeley (2012, p. 9) echoes this figure: “A typical memory span is around six or seven digits, not because the digits themselves are forgotten, but rather because their order is lost.”

But there is a large body of literature that suggests that a figure of 3–4 items is the true capacity of WM. For example, Baddeley (2010, p. R138) puts the capacity of the episodic buffer, the main workspace of WM, at “about four chunks or episodes,” something Cowan (2001) argued for earlier. And others have recently concluded that different models of cognitive capacity “converge on the idea that observers can store around three or four items in working memory” (Cohen et al., 2016b, p. 325).

²²⁴ Such capacity limits on the content of attention, consciousness, or WM are generally measured in “items” of some kind (Chun, 2011, p. 1407). There are other ways of measuring capacities, e.g., area in square degrees of a visual field, but for our purposes, it will be enough to just focus on the items measure of capacity.

In this section I survey three aspects of WM capacity limitations: varieties of capacities; the relation between the number and complexity of items in WM; and that between storage and manipulation capacities. I then survey theories of the underlying mechanisms of these limitations before an interim application to the cases of Florence and Ben.

8.2.1 Varieties of Capacity

There is almost certainly not a *single* “WM capacity,” as should shortly become apparent. On both BH and Cowan’s models, an awful lot of information is being stored and manipulated “in the background,” only a portion of which enters the episodic buffer or the focus of attention and is therefore Executively Attended. There is a sense, then, in which the capacity of WM is much greater than the three to four items, and it may even be difficult to gauge exactly where (non-executive) WM ends and other cognitive systems begin. But of course, our interest here is in that portion of WM that implements Executive Attention.

Here too there are more complexities to be respected. There is ample experimental evidence for different modalities having different capacities for content (the Fractionation Question, 5.3.1.1). For example, Baddeley (2007, pp. 198–199) suggests that while there is good evidence for an overall, general WM capacity, there is also good evidence for more specific or specialised WM capacities. This applies not only to the “slave” components of WM, but also to the central executive, which may be capacity limited either by the limitations of its inputs from the components, or its own capacity for manipulating different kinds of content, or both. In humans, visual content capacity seems to be significantly richer than, say, olfactory content, and certainly richer than interoceptive content. According to Logie and Cowan (2015, p. 317), there is growing evidence for domain-specific WM components, plus domain-general resources that kick in when the former exceed their capacity (called by others, peripheral and central components of WM). The domain-specific

components do not interfere with each other, but there is a drop in performance in specific modalities when domain-general resources are diverted to other modalities. All these different capacities are picked out and measured by different experimental paradigms.²²⁵

Baddeley (2007) speculates on the possible reasons for the finding that “verbal short-term memory is much more readily separable from executive processes than visuospatial” (p. 203). His analysis is roughly that it is easier to rehearse verbal traces and thus maintain them, because they involve over-learned items (like numerical digits) whereas visuospatial traces do not. Thus, more *attention* is required to rehearse visuospatial traces, a resource that is limited in capacity. It is not clear exactly what he means by attention here—is it the Liberal Attention inherent specifically in the middle level visuospatial storage and manipulation, or is it the Executive Attention of the central executive? Indeed, some have considered attention to be the domain-general cognitive resource that imposes the global capacity limits on all cognition, or on WM specifically (Baddeley, 2007, p. 190; Miyake et al., 2000, pp. 88–89). I return to this issue below (8.3).

8.2.2 Number and Complexity of Items

If the capacity of WM is measured in items, does it matter how complex the items are? For example, there is some evidence that a subject can recall fewer long words than short words, an effect that depends both on the number of syllables in the words, and the complexity of the syllables, rather than just their brute duration (Service, 1998). Controversy continues over whether “slot” models or “resource” models are preferable (Eriksson, Vogel, Lansner, Bergström, & Nyberg, 2015, p.

²²⁵ On the intricate interplay between domain-general and domain-specific WM capacities, see also Miyake et al., (2000), Maehara & Saito (2007), Li et al., (2014), and Jarrold et al., (2011).

33,35; Gross, 2018, p. 3; Ma, Husain, & Bays, 2014). On slot models, WM has a limited number of discrete slots for storage of items and their manipulation, while on resource models, storage and manipulation are analogue, continuous resources that can be shared among content in much more flexible ways.

In visual WM capacity Awh et al., (2007) found that visual WM represents a fixed number of items regardless of the complexity of those items, but they did also conclude that the number of representations and the resolution of representations are two dissociable “dimensions” of WM capacity. They even suggest two putative neural loci for the two dimensions (p. 627). A series of seminal experiments (Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001) showed that “objects defined by a conjunction of four features can be retained in working memory just as well as single-feature objects, allowing sixteen individual features to be retained when distributed across four objects. Thus, the capacity of visual working memory must be understood in terms of integrated objects rather than individual features” (Luck & Vogel, 1997, p. 279).²²⁶ Yet Brady and Alvarez (2015) have argued that when items are quite complex, the capacity of WM may be reduced to only 1–2 items, which suggests some kind of trade-off in capacity between the number of items that can be held in WM and their complexity.²²⁷

Nonetheless, Bays et al., (2011) produced results that are difficult to reconcile with a popular model of WM storage on which bound objects are stored and retrieved. They showed that subjects exhibit a pattern of reporting errors in which accuracy for features such as colour and orientation dissociate—that is, subjects may correctly recall the colour, but be mistaken about the orientation of the same object. They take this to support a “shared resource” model of WM capacity, where the central resource cannot be exceeded across discrete feature channels. Hardman and

²²⁶ See also Chun (2011) for a more detailed summary of how different capacities within WM interact with each other.

²²⁷ See also Ma et al., (2014), who conclude that “many details in this framework continue to be debated, particularly the extent to which resources are divisible and the degree to which different features tap independent resource pools” (p. 355).

Cowan (2015) however, conclude that both feature load and object load must be invoked to account for the empirical patterns of reporting failure.

Most pertinently, De Loof et al., (2015) found that a distracting *executive* WM load impacted on visual awareness differently to a distracting *visuospatial* WM load.²²⁸ They hypothesise that the executive WM load raises the threshold of visual awareness for masked stimuli by drawing resources away from “manipulating information” and “decisional processes” (p. 2527)—i.e., by diminishing Executive Attentional further processing that *reports* on the phenomenal visual content, although they admit this hypothesis awaits an appropriate experimental paradigm to confirm it. *A fortiori*, distracting executive load not only hinders awareness, but also *confidence* in what the subject has phenomenally experienced (Maniscalco & Lau, 2015), and a hierarchical model of processing, in which confidence ratings are generated by higher-order processes, were the best fit for the empirical patterns observed (Maniscalco & Lau, 2016).

8.2.3 Storage v Manipulation—Shared resources?

In 5.4 I considered the dual functions of WM—storage and manipulation. Should WM capacity limitations be ascribed to either, or both? It seems the two are intimately connected.

“In this article we use the term WM capacity in a descriptive sense, referring to the fact that people can hold only a limited amount of

²²⁸ De Loof et al., were not in fact measuring effects on consciousness *as such*, but on *report* about conscious content (Chapter 7). Regardless, it is the *difference* in effect between visuospatial and executive WM that is pertinent here.

mental content available for processing. The capacity limit is usually operationalized as a limit on how much new information people can remember over short periods of time (in the order of seconds), but there are reasons to believe (discussed below) that the capacity limit also applies to people's ability to make information in the current environment simultaneously available for processing" (Oberauer et al., 2016, p. 758).

One interesting question is whether storage and manipulation are independent resources with independent capacities, or whether taxing the resources of the one diminishes the capacity of the other. Baddeley (2007) observes that "Bayliss et al. (2003) conclude that their results are inconsistent with an interpretation of complex working memory span in terms of a single resource pool that is shared between storage and processing" (p. 202), yet Cowan (2005, pp. 52–66) reviews the evidence and concludes that storage and manipulation share a common resource at least to some extent, and that that resource is controlled attention. Both these conclusions are consistent with a picture where what can be manipulated by WM is only that which is stored in WM, although not all that is stored in WM is necessarily manipulated, especially by the executive.²²⁹ Efficiency of storage (thus increasing capacity) may be improved by strategies such as grouping or chunking (Brady, Konkle, & Alvarez, 2009), ensemble representations (Brady & Alvarez, 2015), and prioritisation (Myers, Stokes, & Nobre, 2017).

²²⁹ For finer detail on the relationship between storage and manipulation, see Maehara and Saito (2007).

8.2.4 The Mechanism of Working Memory Capacity Limitations

A number of theories have been proposed to explain the mechanisms underlying the capacity limitations of WM described above. These are generally grouped under three categories: the speed hypothesis; the resource pool hypothesis; and the inhibition hypothesis.²³⁰ The underlying concept behind all three is that some content is somehow privileged over other content, and access to the privileged content is maintained for longer, and thereby accessed more by the executive, i.e., Executive Attention.

The *speed hypothesis* posits that WM capacity limits are imposed by time constraints. WM traces decay rapidly over time. Only those that are rehearsed are maintained,²³¹ and there is a limit to how much we can rehearse at any given moment. Those that miss out are lost (Souza, Rerko, & Oberauer, 2014), and thus, only that small number of rehearsed items is maintained in WM, and therefore available for manipulation.²³² The *resource pool hypothesis* posits that WM manipulation is a limited resource that must be doled out to possible objects of manipulation. It may be doled out in a small number of discrete quanta (“slots”) or as a more continuous, flexible resource (8.2.2). WM capacity limitations arise from the limited nature of this manipulation resource. The *inhibition hypothesis* posits that neither the speed of decay nor resource limitations are in play in WM capacity, but that representations compete against each other, until the stronger eliminate the weaker. The victors go on to drink the heady mead of WM manipulation. Thus, on the inhibition hypothesis, WM capacity limits arise ultimately from the elimination of all competing representations except a few.

²³⁰ I have used terminology from Baddeley (2007), Chapter 11, although other authors use slightly different terms. For example, Oberauer et al., (2016) call the three approaches: temporal decay (=speed); limitation in cognitive resources (=resource pool); and mutual interference (=inhibition). These kinds of mechanisms have also been expressed in terms of predictive coding (Hohwy, 2012, pp. 6–7). Rhodes and Cowan (2018) argue for an additional mechanism that involves interactions between WM and LTM.

²³¹ This mechanism has been challenged (Lewandowsky & Oberauer, 2015).

²³² For a recent review of the speed hypothesis, see Ricker et al., (2016).

Oberauer et al., (2016, p. 759) explore the prospects of a parsimonious single factor explanation of WM capacity, where just one of the three hypotheses is solely responsible for WM capacity limits. But this quest is unlikely to succeed: “it is unlikely that working memory span, with its capacity to predict performance on a rich array of cognitive tasks, can be adequately interpreted in terms of a single variable” (Baddeley, 2007, p. 198). In similar vein: “the existing data do not appear to universally support any one of the three accounts of working memory capacity” (Oberauer et al., 2016, p. 758). This might be at least partially because the three hypothesis overlap each other conceptually. For example, *rehearsal* capacity plays a central role in the speed hypothesis, but it is a function of the limited *resources* available to rehearse multiple contents simultaneously, which strongly resembles the resource hypothesis.²³³

In all three hypotheses, attentional strategies are clearly implemented—rehearsal (abundance of access / maintaining access), influence on posterior processing, filtering, exclusion, and competition. Attention is intimately embedded in the manipulations of WM and plays a vital role in the mechanisms of WM capacity limitations. What is more, the *biased competition* model of attention (Desimone & Duncan, 1995; Ruff, 2011) described in 3.3.4.2.3—which emphasises that competition is not free and fair but biased by influences from memory, subconscious heuristics, and so on—expands on the inhibition hypothesis of WM by highlighting that WM is inseparable from the rest of cognition, and that attention is an important bridge.

²³³ See also Oberauer and Lewandowsky (2011) for a detailed model of how these two hypotheses may interact.

8.2.5 Back to Florence and Ben

The “capacity of WM” is no simple matter. The selective survey above has nonetheless served to highlight a number of important points that strengthen the arguments in the previous chapter and serve to underwrite those that follow in this. Appreciating the varieties of capacity in WM (8.2.1) indicates that there is a great deal of content in WM that is not Executively Attended, and indeed, it is difficult to draw clear boundaries between WM and the rest of cognition.²³⁴

In considering the relationship between capacities for the number and complexity of items (8.2.2), we saw there is good evidence that WM stores whole, bound objects, although the precision or fidelity of that content is limited (much less than the dozens of speckles on Ben’s hen) and degrades with increasing numbers of items simultaneously stored. De Loof et al., suggested that it is Executive Attention that is responsible for this degradation. That is, Executive Attention in WM is what limits what can be reported by the subject, including the degree of fine detail in an item. Its effects dissociate from the effects of non-executive loads, such as visuospatial ones.

What is more, there is an effect on confidence that matches the effect Florence experiences when she reflects on her peripheral content and her confidence drops. But as predicted, Ben’s single item foveal experience of the hen shows no such diminution in confidence. The hierarchical model espoused by Maniscalco & Lau accommodates this difference and matches the Witches’ Hat Model I describe below (8.5). Further still, prioritisation and chunking—strategies employed by the WM executive—do not describe Ben’s phenomenal FOC of the speckles. Rather, he sees them as a whole, simply, as they are (or appear). The only chunking is the unity of his experience as a whole.

In 8.2.3 we saw that not all that is stored in WM is manipulated by the executive. Overflow, as I have framed it, requires only consciousness without Executive Attention. That could plausibly describe such unmanipulated WM content, if the

²³⁴ See also 5.5 and 8.5.

mechanism for content to become phenomenal need not involve Executive Attention (*contra* Phenomenal definitions of attention, 3.3.3). I argue below for the stronger conclusion that content completely outside WM can become phenomenal, but even un-Executively Attended content within WM would still be enough to establish the $C \sim A$ of overflow. Whether Ben's rich speckle content is within or outside WM is not what matters here—only whether it is Executively Attended.

Finally, it is clear that all the hypothesised mechanisms underlying WM capacity (8.2.4) implement attentional strategies. These strategies subserve higher-order processing, such as Experience-of and Seeming-that, but are not intrinsically inherent in the production of phenomenal FOC as such. None of the four characteristics of phenomenality (2.4) require them.

Clearly, Executive Attention is inextricable from WM capacity limitations, yet some have studied attentional capacity limitations in their own right, sometimes coming to conclusions that contrast with those above. A brief survey of that body of knowledge is also in order.

8.3 The Capacity of Attention

As I observed in 5.4.2, Executive Attention has not uncommonly been identified with the manipulation aspect of the executive of WM. Yet the literature on the capacity limitations of attention diverges in some ways from the patterns observed in the literature on WM capacity limitations. In this section, I consider what the concept of capacity might mean when applied to my definition of attention as the structuring of cognition for further processing, and then outline some pertinent empirical findings about the nature and magnitude of this capacity. I conclude by considering how attentional capacity illuminates the cases of Florence and Ben.

8.3.1 The Nature of the Capacity of Attention

On my definition of attention as a suite of *strategies* (Chapter 3), there seems no principled reason to think of it as a limited-capacity resource within a cognitive economy. That would be like thinking of “running” as a limited-capacity resource at the Olympic Games: the more races are held, the less running is left over, because all the running has been used up by the athletes. That’s not how it works. The amount of running that can occur depends on the number of people available who can run, and on how far they can run, but not on the “running” itself. It is the specifications of the mechanism—the resources available to implement attentional strategies—that limit the amount of attending that can happen in a cognitive economy, not the nature of attending itself.

The resources in question are, if we identify Executive Attention with the cognitive executive, just the resources of the cognitive executive. We saw in 8.2.4 above that the hypothesised mechanisms of WM capacity rely upon attentional strategies. That being so, one would expect the empirical pattern of Executive Attentional capacity limitations to be identical to those of WM. But as we shall see in the next section, there are some interesting contrasts.

A possible reason for this contrast is that not all researchers make such a close identification between Executive Attention and the cognitive executive. One view is that attention is the gateway to WM (Prinz, 2012), the path by which content enters WM and becomes available for manipulation. This view does not deny the role of attention within WM manipulation, but extends its role to pre-WM selection. On this kind of view, attention and WM could indeed exhibit different capacity patterns—for example, more content might be encoded into WM by attention than WM can manipulate. However, on the definitions I am using here, this pre-WM attention would constitute Liberal rather than Executive Attention, what Wolfe (1994, p. 202) calls “a preattentive, massively parallel stage that processes information about basic visual features” and distinguishes from “a subsequent limited-capacity stage that

performs other, more complex operations.” My focus in this section is on the latter kind of attention.

Another possible reason for the contrast is the difficulty in operationalising measures of attentional capacity.

“Although complex processes often cannot be reduced to a single number, summary indices like capacity for working memory and *gF* for fluid intelligence are useful for quantifying individual differences and changes in abilities over time. A comparable measure of attention – an objective, standardized summary score – would benefit both research and clinical practice” (Rosenberg et al., 2017, p. 291).

The authors go on to highlight some of the difficulties inherent in any attempt to measure a single capacity limit of attention (pp. 299-300, boxes 1 and 2). Chief among these is the fact that attention is not a unitary entity, but a suite of quite diverse entities operating on diverse classes of targets. This means that any given subject’s attention capacity will vary with the kind of attention being employed, and even over time, as the subject performs different tasks in different situations:

“For example, one person may struggle to pay attention for long periods of time but have no difficulty switching between tasks whereas another may have no trouble maintaining uninterrupted focus but lack the ability to multitask” (p. 299).

Like WM, the “capacity of attention” as it is commonly discussed in the literature is better thought of, then, as a suite of individual and specific capacities for different attentional tasks (e.g., search, tracking, orienting, etc.) in different modalities (e.g.,

visual, auditory, etc.) that varies from individual to individual and even within the same individual over time. What is more, even within a single modality, such as visual attention, it will be complicated by the different kinds of attention possible, in this case, object, feature, or spatial attention (3.2.3). An overall “capacity of attention” should perhaps be viewed with the same suspicion as an overall “intelligence quotient” (IQ), and for similar reasons—because they attempt to sweep up a large number of intricately interacting cognitive processes with different capacity limitations in a single net. This fits in nicely with my approach of taking attention to be a suite of strategies implemented by a range of systems and subsystems whose capacity limits depend on the specifications of each system, and interact in complex ways. But there is no reason to suppose that attentional strategies (which are conceptually derived) map neatly onto WM components or subprocesses (which are empirically derived). If so, differences in capacity patterns between attention and WM ought not surprise us.

8.3.2 The Capacity Limitations of Executive Attention

Nonetheless, people do speak of the overall capacity limits of (Executive) attention. For example:

“In terms of visual attention, initial studies estimated that around three or four locations can be attended at once, but more recent efforts have pushed that number closer to around seven or eight” (Cohen et al., 2016b, p. 325).

Franconeri et al., (2007) point out that up to the early 1990’s, visual spatial attention was thought to be unitary, focusing on one spatial location in the visual field at any one time. Subsequently, a number of studies showed that subjects could be cued to

spatially attend up to five locations simultaneously, as evidenced by the improvement in their recognition of stimuli when cued, as opposed to uncued stimuli. They then present their own evidence for spatial attention to up to eight simultaneous locations. Interestingly, however, there does appear to be a trade-off between the maximum *number* of locations that can be attended and the *precision*²³⁵ with which a location needs to be attended to—higher precision means fewer locations. They ponder whether the relatively high precision required in previous experiments might have been the reason for the observed limits of one or up to five spatial locations simultaneously. In other words, their evidence suggests that limitations on the number of spatial locations that can be simultaneously attended do *not* arise because “the visual system is somehow architecturally restricted to dealing with a fixed number of items at once” (p. 1004) but because of the trade-off between precision and number of locations. Similar findings to these in spatial attention have been found in relation to rapid enumeration, visual search, and multiple object tracking (p.1011).

These findings suggest that when Executively Attending to multiple spatial locations or tracking objects, the standard WM capacity limitation of 3-4 items often does not apply. Executive Attention capacity does not neatly map onto WM executive capacity, at least on the empirical paradigms discussed above. But how far can this go? Is there any limit to the capacity of divided spatial attention? More importantly, can attention be divided so far as to attend to every speckle on Ben’s hen simultaneously?

²³⁵ The authors manipulated precision by varying the number of locations within the test field. A crowded field requires more precision to identify the correct spatial cues than a sparse field (p. 1005). See also Wu and Wolfe (2018) and Cohen’s (2019) response.

8.3.3 Back to Florence and Ben

There is a limit to how far Franconeri et al's trade-off between number and precision can be taken.

“This effect of this interelement crowding seems to reflect a lower limit on the precision of selection when a single item is selected. When the distance among elements is denser than this level of precision, elements can no longer be selected independently (Intriligator & Cavanagh, 2001)” (Franconeri et al., 2007, p. 1010).

There is a limit to the precision with which spatial locations can be selected:

“Our data show that selection has a coarse grain, much coarser than visual resolution ... The results suggest that the parietal area is the most likely locus of this selection mechanism and that it acts by pointing to the spatial coordinates (or cortical coordinates) of items of interest rather than by holding a representation of the items themselves” (Intriligator & Cavanagh, 2001, p. 171).

Executive Attentional selection likely operates via a different mechanism to that of visual experience.²³⁶ That is hardly a surprise—holding content in one's experience is one thing, selecting items among that content is clearly another. But the interesting thing here is that “selection has a coarse grain, much coarser than visual

²³⁶ Usher et al., (2018) argue from the ability to discriminate targets outside the focus of attention to phenomenal overflow. That is not the argument I am making here. On my definitions, the ability to discriminate is intrinsically attentional, involving higher-level processing. Further, it can be performed implicitly, without consciousness.

resolution.” For Florence, peripheral visual resolution is coarse because of the specifications of the visual system. It may be (though it cannot be assumed) that her peripheral phenomenal visual resolution is roughly as coarse as her peripheral spatial resolution.

Not so for Ben. The speckles on the hen are significantly denser than this threshold of attentional selection. Ben cannot divide his attention among them all, both because of their density and their number (many dozens). Clearly, Ben *cannot* attend to every speckle simultaneously. While there are some differences in the capacity patterns of Executive Attention and WM (possibly for reasons explained in 8.3.1), both patterns share this tight constraint on their resolution. The question that remains, then, is whether Ben’s phenomenal visual resolution is so constrained.

8.4 The Capacity of Consciousness

In many ways the capacity limitations of consciousness are more difficult to handle than those of either WM or attention. This is due in no small part to the enigmatic nature of phenomenal consciousness itself, in contrast to the relatively more easily defined operational characterisations of attention and models of WM. Nonetheless, the contrast between the capacity for conscious content and that for attended or WM content is significant, and counts as a strong reason to think that there is indeed $C \sim A$.

In this section I argue that as defined, phenomenal consciousness has no intrinsic capacity limitations. The only limit to phenomenal content lies in whatever systems are responsible for producing that content. In humans, these systems are causally complex, and the burden of proof is upon the overflow sceptic to show that only executive WM content can be phenomenal, and that non-executive WM content is

never phenomenal. But this is implausible for a number of reasons, among which are the dissociation between the neural substrates of attention and consciousness, and the indubitable richness of non-WM stores such as iconic memory. Given these factors, it would be very remarkable indeed for the capacity limitations of consciousness not to differ from those of attention.

8.4.1 The Features of Consciousness

In 2.4 I defined phenomenal consciousness using four features: what-it-is-likeness; a situated first-person perspective; phenomenal unity; and temporality. I further made a distinction between the content of phenomenal experience and phenomenality as such. If there is a capacity limitation to conscious experience, it seems likely that such limitation would be a limitation on the amount of content (perhaps the resolution or grain of that content) rather than a limitation on phenomenality itself.²³⁷ The Target question, which is central to our identification of instances of $C \sim A$, applies not to phenomenality, but to content—is this Target simultaneously conscious but not attended? $C \sim A$ is not a matter of clarity or vividness, is not a matter of being more conscious or less attentive. To answer Q, ideally we need clear cases of phenomenal consciousness of a Target which is not in any way Executively Attended.

If one accepts phenomenal unity, a situated subjective perspective, and temporality as characteristics of consciousness, then in one sense, the capacity limitation of

²³⁷ Physiological consciousness—the kind measured by the Glasgow Coma Scale (Reith, Brande, Synnot, Gruen, & Maas, 2016) or the Bispectral Index of sedation (De Deyne et al., 1998)—is undoubtedly graded. But gradations in *what it is like-ness* itself—whether globally for a subject, or particularly in the experience of specific targets—is a more contentious issue (Bayne et al., 2016; Morin, 2006; Overgaard, Rote, Mouridsen, & Ramsøy, 2006; Sergent & Dehaene, 2004; Windey & Cleeremans, 2015). While fascinating, it is not directly relevant to my arguments here.

phenomenal experience is just *one*. A subject can only have one unified experience at a given moment in time.²³⁸ But of course, what is pertinent to the Capacity Question is whether there is any kind of limitation on the “quantity” of the multiple *contents* of a unified phenomenal experience.

Like attention-as-strategies (8.3.1) there is no limit on the amount of content that can become phenomenal. It is plausible that humans experience far more phenomenal content than goldfish. Dogs experience far more phenomenal olfactory content than humans. There is nothing inherent in consciousness itself that dictates a limited capacity. Actual limitations are imposed by the systems that produce the content. The crucial question is, therefore, *which* systems can contribute to phenomenal content, and can they do so without employing Executive Attention?

There are good reasons to think that phenomenal content has a far finer resolution than Executively Attended content. According to Raffman (1993), “whereas we can experience and distinguish 1,400 frequencies in sound (through same/different judgments), we can report on and identify only 80 pitches, purportedly showing that phenomenal consciousness has the capacity of around 1,400 pitches whereas access consciousness has the capacity of only around 80” (Jennings, 2015, pp. 284–285). The suggestion here is that a subject is capable of being phenomenally conscious of far more nuance than she can explicitly report through WM channels. Caution is required in drawing conclusions from this result—e.g., Raffman is talking about discriminatory capacities over time, whereas in this chapter I am interested in synchronic capacities. But his results do suggest that the grain of auditory phenomenal FOC (as measured by 1,400 frequency distinctions—Seemings-that two sounds are different) is much finer than the grain of Experience-of report (as measured by 80 pitch identifications).

Another way to make sense of this is to go back to my definitions of attention, consciousness, and WM, and consider the type of relationship that connects them

²³⁸ Thus also Baars (1997b, p. 368).

(4.4). On my view, the executive of WM is the mechanism that implements Executive Attentional strategies to manipulate content, and phenomenal consciousness is an attribute of some or all of the output of this mechanism. The relationship that best describes this account is a complex causal chain (subset of CA). Implementation of attentional strategies is one of the causes of content becoming phenomenal ($A \rightarrow C$).²³⁹ There are very likely at least sometimes intermediate steps ($A \rightarrow X, Y, \text{etc.} \rightarrow C$). It is also likely that becoming conscious of certain content can attract bottom-up attention to it ($C \rightarrow A$), and so on. Overflow scepticism imposes a constraint on these complex causal chains: it denies the possibility of $X, Y, \text{etc.} \rightarrow C$ without Executive Attention being one of the causes of consciousness. Thus, the capacity of consciousness can never exceed the capacity of the Executive Attention that causes it, since only content that is Executively Attended can become conscious.

This constraint, and the transmission of capacity limitations from Executive Attention to consciousness, seems implausible for at least three reasons, two of which I consider briefly here and the third in more detail in the next section. First, this kind of complex causal chain means many other causes apart from Executive Attention are involved in producing the output of conscious content. Whether these other causes are capable of producing that output *without* Executive Attention is something to be established, not assumed. Keep in mind that *Liberal* Attentional strategies are almost certainly necessary, but the necessity of *Executive* Attention is an open question.

Second, precisely which cognitive apparatuses (a) constitute the cognitive executive and (b) are responsible for consciousness, remain open questions. The overflow sceptic requires either both to be identical, or for the apparatus of Executive Attention to subsume the apparatus responsible for consciousness. We don't yet know enough about either to assume either proposition.

²³⁹ The arrows signify the concept of causation generally without engaging in the deep waters of its metaphysics.

8.4.2 Neurophysiological Considerations

This raises the third objection to the sceptical constraint above—the Neural Question (5.3.2.6). There is little support from our current picture of the *neural substrates* of Executive Attention and phenomenal consciousness for the overflow sceptic’s account. My impression is that the literature on the degree of coincidence of the neural substrates of Executive Attention and consciousness is inconclusive at this time. Nonetheless, there are clues: the strong evidence for a significant difference between the neural substrates of Executive Attention and consciousness casts serious doubt on the constraint; and the phenomenon of neural divergence is good reason to accept the richness of non-WM stores such as iconic memory, and therefore of FOC (and perhaps, therefore, of phenomenal FOC). I consider each in turn.

8.4.2.1 *Different Neural Substrates*

It has been suggested that attention and consciousness share the same neural substrate (O’Regan & Noë, 2001). But at least two recent reviews of the evidence relating the neural substrates of attention and consciousness conclude that they significantly dissociate neurally (Lamme, 2004; Tallon-Baudry, 2012), a conclusion shared by others (e.g., Baars, 1997b, pp. 367–368; Kanai et al., 2006; Koch & Tsuchiya, 2012). The neural *markers* of Executive Attention and consciousness also doubly dissociate (Koivisto & Revonsuo, 2007; Wyart, Dehaene, & Tallon-Baudry, 2012), and a case can be made for them evolving along different paths (Montemayor & Haladjian, 2015, Chapters 2.2, 5). Even the ensemble statistics invoked by overflow sceptics to explain away the “illusion of richness”—which are plausibly Experience-of-type content—seem to have neural pathways distinct from the rest of perception (Cohen et al., 2016b, pp. 325, 330). This builds a powerful case for a significant dissociation between the neural substrates of Executive Attention and

consciousness (always keeping in mind uncertainty about either substrate), and therefore undermines the transmission of capacity limitations from the one to the other.

8.4.2.2 *Iconic Memory*

It has been fairly well established for some time now that there is at least one fragile short-term memory store outside of WM—iconic memory (G. M. Long, 1980).²⁴⁰

Iconic memory is a prominent candidate for an example of this dissociation of the neural substrates of Executive Attention and phenomenal consciousness. There is evidence that the content storage capacity of iconic memory far exceeds the content manipulation capacity of Executive Attention.

Iconic memory differs from WM in two chief aspects: (a) it seems to contain much richer content (phenomenal or not) than WM; and (b) it degrades more quickly than WM. But it interacts with WM such that a small proportion of the content in iconic memory can be “transferred” (whatever that may mean) to WM—perhaps through mechanisms such as rehearsal or competition—where it may be used for more complex processing. Iconic memory beautifully explains the Sperling (1960) experiments. Recall that here, a subject is exposed briefly to a grid of 12 or 16 letters, and subsequently cued to recall a row of them. Subjects are generally reliable in reporting the cued row, but are then unable to report any other rows. This holds regardless of which row is cued, suggesting that all the content—every single letter—was encoded into iconic memory, but once a subset of those letters—the cued row—was transferred (whatever that means) to WM for report, the traces of the rest of the letters degraded and were lost. This suggests that the content of

²⁴⁰ See also Landman (2003, p. 162), Vandenbroucke et al., (2011). Similar arguments may be mounted for other short-term sensory memory stores, such as auditory echoic memory (Zimmermann et al., 2016).

iconic memory—the whole grid—is richer than the content of WM—just one row—which is (a) from above; although the persistence of the iconic memory traces is much shorter than the persistence of the WM traces, which is (b), (Bronfman et al., 2014, p. 1394).

Consideration of visual system neurophysiology (Schwartz, 2017, Chapters 2, 3) provides further reason to think that iconic memory content may indeed be rich beyond the ability to report. Here, we find a pattern of initial convergence and shrinking of capacity, followed by a much greater divergence and expanding of capacity. Over 100,000,000 photoreceptors (120,000,000 rods and 6,000,000 cones, p. 37) converge upon just 1,000,000 ganglion cells. There is further convergence through the lateral geniculate nucleus (LGN) to the striate cortex (V1). But from there, there is a significant divergence, spreading out from V1 first to the extrastriate cortex which is specialised for analysing attributes such as motion and colour, and then “to higher centres, which combine visual information with memory and other senses. Higher visual centres, in turn, send information back to the striate cortex (V1) via reciprocal projections” (p. 20). Thus, from the striate cortex on, the richness of capacity and processing is *increasing*, in contrast to the *decrease* in capacity from the lens through to the striate cortex.²⁴¹ There remain open questions as to whether this rich content is attended and/or conscious, but neurophysiological considerations certainly give us good reason to think iconic memory stores to be *rich*.

²⁴¹ Note that my FOC and higher-order content don’t map perfectly onto this neural geography. Higher cortical structures have copious feedback connections to lower ones (Montaser-Kouhsari & Rajimehr, 2004, p. 434), as far down as the LGN. Determining the actual neural correlates of my threefold distinction of FOC, Experience-of and Seeming-that is a complex topic that unfortunately I cannot explore here. See also Lamme (2018, p. 2) and Haun et al., (2017) for retinal arguments about Florence cases.

8.4.3 Back to Florence and Ben

To summarise the above (8.4.1 and 8.4.2), neither Executive Attention-as-strategies or phenomenal consciousness as I have defined them are intrinsically capacity limited. Rather, it is the storage and manipulation capacities of systems like WM or iconic memory that impose the limitations. Not only do Executive Attention and phenomenal consciousness seem to have dissociable neural substrates, but consideration of the neurophysiology of vision boosts the plausibility of iconic memory stores rich beyond the manipulation capacity of Executive Attention in WM. We thus have strong reasons to hold that the capacity limitations of Executive Attention and phenomenal consciousness are not identical.

Ben's iconic memory holds rich foveal speckle content while Florence's may not hold rich peripheral tree and leaf content. The critical question is whether Ben's overflowing iconic memory content is both phenomenal and unattended. The pattern of trading off precision for number that characterises Executive Attention in WM (8.3.3 above) does not apply to phenomenal consciousness. Ben's foveal phenomenal vision, underwritten by the richness of iconic memory content and the neurophysiology of the visual system, makes no such sacrifices in resolution. Florence's peripheral phenomenal vision does make such a trade-off and its resolution is severely limited by the neurophysiology of the visual system. She has knowledge of her trees and leaves, albeit imperfect (Cohen, 2019; C.-C. Wu & Wolfe, 2018). But Ben has at least perfect phenomenal experience of every speckle, even if his Executive Attention-dependent knowledge (Experience-of, Seeming-that) is limited and imperfect. So long as he can visually resolve a hundred speckles, he can phenomenally experience—as richly as his visual system allows—every speckle. Largely peripheral phenomena such as crowding (Pelli, 2008; Whitney & Levi, 2011) while pertinent to Florence, simply do not apply to Ben.²⁴² When he tries to spatially attend, though, resolution of attention degrades far more quickly (8.3.2).

²⁴² E.g., “when the same display appears at the fovea, the oriented items are not crowded and the orientation signals do not appear to be obligatorily averaged” (Alvarez, 2011, p. 172).

The critical question here is just *how* it is that content becomes conscious. If we knew that, we could work out whether the rich content in iconic memory does indeed become phenomenal independent of Executive Attention in WM. But unfortunately, we don't know that, and there are major obstacles to our ever knowing that (7.4).

That is why this is not a deductive argument for overflow, but an abductive one. So far, I have marshalled the evidence for a model where capacity limitations of attention and consciousness dissociate. Given this, it would only be the most remarkable of coincidences if empirically those capacity limitations were found to be consistently identical. What is more, there is good reason to think that FOC in, say, iconic memory is quite rich, far beyond the capacity of Executive Attention in WM, and uncertainty about how it is that content becomes phenomenal. Taken together with my arguments for the veridical immediacy of Ben's rich speckle experience in Chapter 7, the best explanation is one that allows for phenomenal overflow.

In the balance of this chapter, I bring together the ideas above into a rough model of how content is processed that provides a convenient way to express them, and then consider the chief alternatives to the overflow interpretation.

8.5 The Witches' Hat Model

“ Rather than attempting to locate consciousness within the box-and-arrow diagrams beloved by cognitive neuroscience, it might be better to think of consciousness as involving a dynamic unity that

is *superimposed* on the relatively static structure of thought and perception” (Bayne, 2010, p. 245).

Memory stores in humans, and their access and manipulation, are not unitary. There is now substantial evidence for different systems with different substrates, functions, capacities, and patterns of behaviour. I have already touched upon the difference between WM and iconic memory above. These two systems display the pattern described by Franconeri et al, (8.3) that is characteristic of attentional processes (but *not* phenomenality) whereby quantity is traded-off for quality. Attentional processing increases refinement of content at the cost of reducing content item capacity. Building on this central idea we can derive what I call a *Witches’ Hat Model* of the flow of content through cognition—a model²⁴³ that captures the features discussed in this chapter thus far.

It is highly likely that WM and iconic memory are not the only short to medium-term visual stores. Crick and Koch (1990, p. 269) and Lamme (2003, p. 14) describe a model with just these two, but Block elaborates further:

“The upshot is that the first phase is very high capacity and is over by 1,000 msecs; the second phase is high capacity and lasts up to 4 seconds; and the third phase has a similar capacity to the working memory phase in Sperling and in Landman et al.” (Block, 2007, p. 491).

Similarly, Sligte et al., (2008; 2010) postulate a “fragile visual short-term memory” (VSTM) with capacity and duration intermediate between brief, rich iconic memory

²⁴³ The WHM might be better described as a metaphor than a model in that it does not seek to simulate as closely as possible the reality being explained, but rather to draw a connection between central or salient features common to both reality and metaphor (5.3.3).

and persistent, sparse visual WM.²⁴⁴ And Haun et al., (2017) sketch a picture of rich short-term memory stores that require time for some of their content to be refined and further processed in different ways.

There are differences in detail, but for my purposes, that is where the devil is *not*. These examples serve to illustrate the general principle captured by my WHM of the gradual trade-off as further processing occurs between the richness and the number of items (Figure 5).

I have depicted three phases²⁴⁵ following Block and Sligte et al., but the boundaries between the phases (and indeed, the boundaries with the rest of cognition²⁴⁶—the external surfaces of the hat) should be taken to be very fuzzy indeed. I am a great believer that brains are organic organs and operate on principles far more like livers and bone marrows than silicon chips or quartz watches. The rest of our physiology is “mushy”²⁴⁷ in this way, with very few distinct functional boundaries, and interacting systems and subsystems melt into each other. It would be surprising if brains were the exception.

²⁴⁴ Phillips (2018, p. 3) discusses some of the challenges that have been raised against a distinct large capacity short-term store.

²⁴⁵ These three phases must not be confused for the threefold kinds of content I described in 7.2.1. While FOC may mostly arise in the brim of the witches’ hat, Experience-of and Seeming-that do not correspond to either of the two higher-order phases of the WHM, but may both cut across them.

²⁴⁶ For example, the boundaries between WM and long-term memory may be quite fuzzy (Brady, Störmer, & Alvarez, 2016; Rhodes & Cowan, 2018). This is the *Integration Question* (5.3.2.4).

²⁴⁷ I am developing the idea of this “Mushiness Principle” further in a future paper.

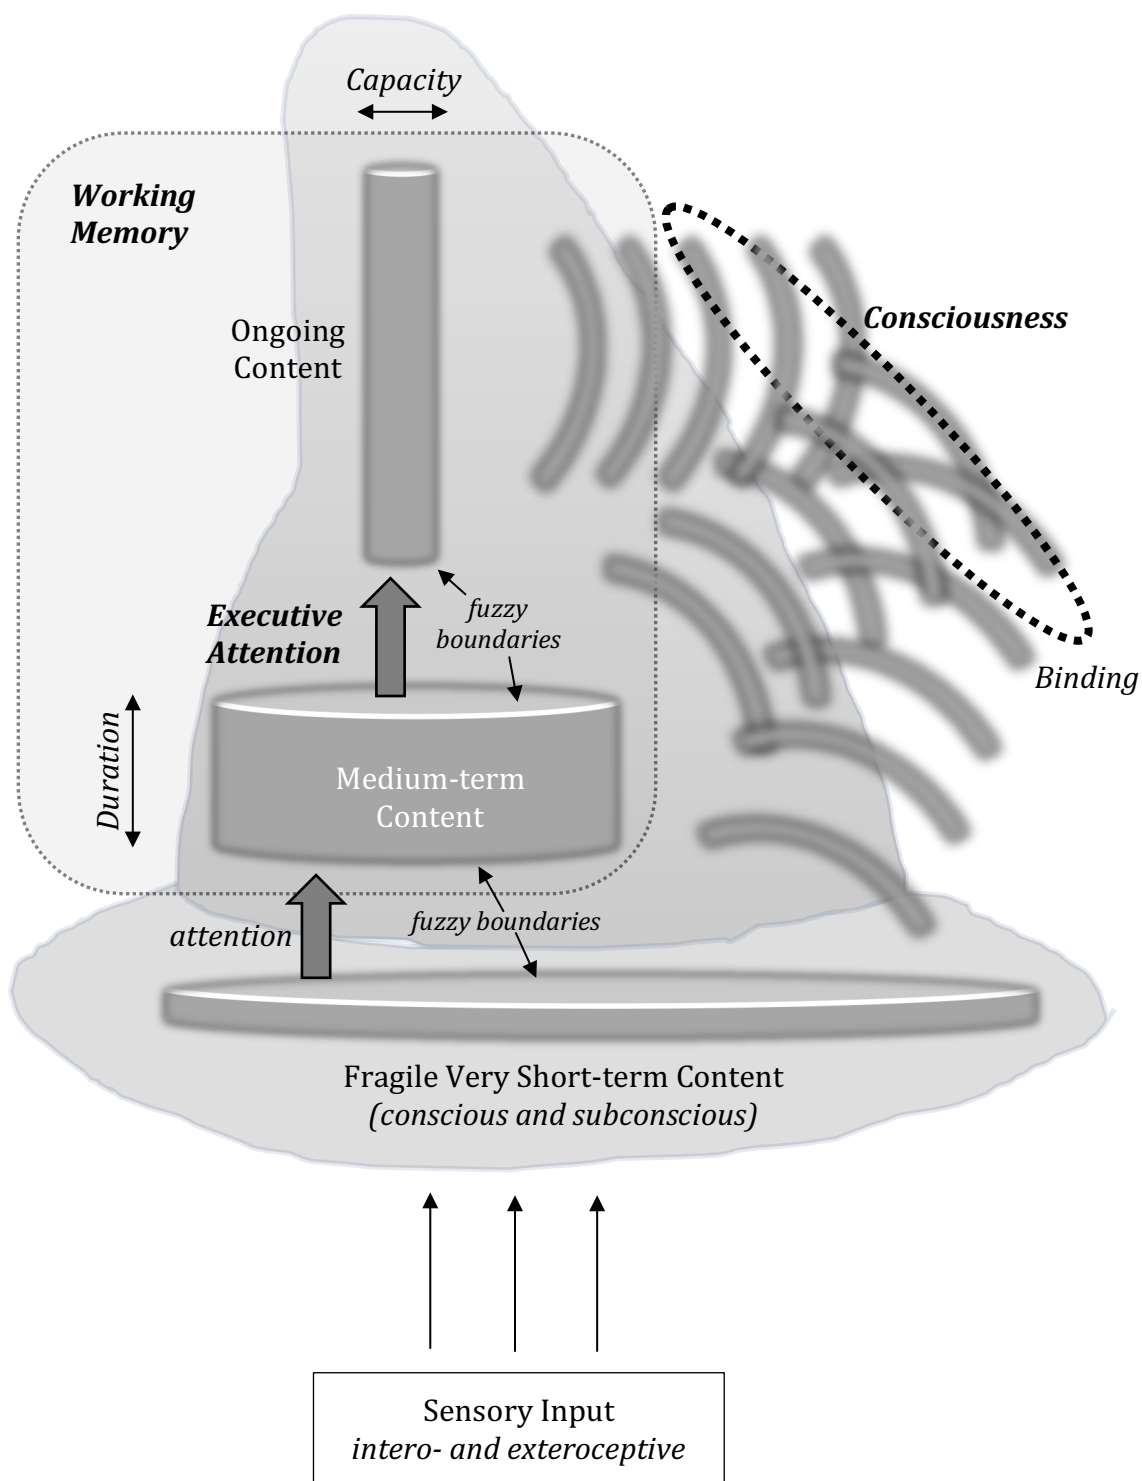


Figure 5 Witches' Hat Model of Conscious Content. I have depicted just perception for simplicity, but the model could be extended to encompass other cognitions: long-term memory encoding and retrieval, imagination, efferent motor command, etc.

In fact, it would be better to think, not of three independent stores with different specifications, but of just one “hat”—a *single store with three overlapping phases*, whose capacity to hold onto items degrades predictably over time, but in different ways. The discrete stores described by Block, Sligte et al., and others are delineated by the experimental paradigms used to investigate them, but there is no reason to think, for example, that content can only persist for either one second or four seconds, but not any period of time in between (Block, 2007, p. 491). Both the persistence durations and the amount of content vary *continuously* among the various stores, rather than being strictly *quantised*, although there may be “bulges” along the continuum.

In the diagram, the horizontal axis represents overall capacity for richness of content while the vertical axis represents the maximum duration over which content can be maintained. The very wide brim tapering up into the narrower peak of the crown thus captures the trade-off described above. The vertical, temporal axis also represents in a general sort of way the idea that content “moves” (whatever that may mean)²⁴⁸ from perceptual apparatus first into the lowest phase, the brim, and thence gradually up through the crown of the hat, or possibly no further if it is nearer the edge of the brim. Perceptual content moves into the brim mostly by being Liberally Attended.²⁴⁹ Movement from the brim into the crown is attended, but depending on how one defines the cognitive executive—and therefore Executive Attention—this may be either Liberal or Executive Attention. The boundaries of WM will be defined with reference to Executive Attention and whether one considers the

²⁴⁸ I make no definite claims about what “moves” signifies here. WHM is neutral with respect to the Manipulation Question and the related Duplication Question (Chapter 5 and 7.2.1.3)—whether separate duplicates of content are encoded in the different phases, or the content remains physically in the same neural substrate but activated differently, somehow, or something else entirely.

²⁴⁹ Note that not all persistence is due to attention. The persistence of a retinal trace, for example, occurs independent of attention (except perhaps orienting?). This is true of the earlier levels of processing, and becomes gradually less true as one moves up the processing levels.

intermediate stores to be part of WM (perhaps the visuospatial sketchpad) or not (Sligte et al's fragile VSTM).²⁵⁰

These border disputes do not impact heavily on Q and are matters for empirical clarification. What is central to the WHM is the idea rich unattended FOC, some of which is further processed by attention to differing degrees, and so moves up the witches' hat. The crucial difference between WHM—which is friendly to the accounts of overflow proponents like Block—and WM models like those of BH and Cowan—which make overflow impossible—is that content *anywhere* in the witches' hat can and regularly does become phenomenal, even that which never moves into the crown. Processes of binding are occurring from the earliest visual processing areas onwards (Holcombe & Cavanagh, 2001; Lamme, 2010, p. 211) and all the way to the peak. Ultimately, what we take to be our ongoing stream of consciousness is composed of FOC (brim) and higher-order content (crown) all bound²⁵¹ and integrated together into a single unitary phenomenal experience.

This model has a number of advantages. First, it explains the immediacy of phenomenal FOC (7.4.2) since content unprocessed by Executive Attention can become phenomenal. It underwrites the confidence Ben has in his immediate phenomenal FOC, since further processing creates more opportunity for errors (Gross, 2018, p. 5), whereas immediate perception leaves less room for error. Models that make Executive Attention necessary for consciousness do not have this explanatory option. Second, it makes sense of the observed ongoing intricate interplay between attention, consciousness, and WM. I do not have the space here, but a great deal of empirical evidence fits neatly into the WHM. Thirdly, unlike the models of BH, Cowan, etc., it does not beg the question of whether consciousness can

²⁵⁰ Trübutschek et al's (2017) description of content held non-consciously in the primary visual cortex and recalled seconds later sounds very much like it is held in iconic memory. What is more, it is "activity-silent" on MEG. However, they describe it as non-conscious WM—an illustration of how difficult it can be to draw the borders. For two possible neurophysiological accounts similar to the WHM, see Magnussen (2009) and Lamme (2010, p. 213).

²⁵¹ The role of binding in the WHM is unfortunately too large a topic to address in detail here.

arise without Executive Attention or WM. WHM allows for Executive Attention-free conscious content, but also leaves the door open for the possibility that all conscious content might actually be Executively Attended—if, for example, it is shown that although it is possible, no brim content *actually* becomes phenomenal, or that Executive Attention reaches down all the way to the edges of the brim of the hat.

8.6 Alternative Interpretations

“In philosophy, there are the wide and narrow roads as well as in religion, and the wide road that leads to destruction involves the smugness of thinking that one can establish one's own view merely by undermining the alternatives. Defenders of a view owe us not only objections to alternatives, but also a worked out version of the view that can withstand scrutiny on its own” (Kvanvig, 2007, p. 167).

I have so far made a case for the WHM on which foveal phenomenal content can overflow attention because it can arise independent of Executive Attention in WM. But I have stressed that mine is an abductive argument, an inference to the best explanation. Having described my “worked out version” of a positive view, it remains then to consider the chief alternative accounts and show why the WHM is a better explanation than they. That does not make the WHM “right,” only preferable. In this section I analyse three such alternatives: the illusion of richness; virtually unlimited Executive Attention; and inchoateness. The accounts I consider here are by no means exhaustive—there are other ways of being either an overflow proponent or sceptic.

8.6.1 Illusion of Richness Account

In this section I compare my overflow-friendly WHM account of Ben's hen to the overflow-sceptical "illusion of richness" account based on ensemble statistics (7.1). There are a number of advantages to the latter, two of which are highly relevant here. First, ensemble statistics allow the capacities of attention and consciousness to be equal without impoverishing either conscious experience or our functionality in the world (Cohen et al., 2016b, pp. 325, 332). Statistical averages of content peripheral to attention are enough for effective functionality, but not so demanding that they tax cognition with irrelevancies.²⁵² What is the point of being richly conscious of content that you never use for any practical purpose? Second, it is parsimonious. The epistemic barrier (7.4) suggests that we can never find conclusive evidence for phenomenal overflow, since we can only identify phenomenal content via report, which requires Executive Attention. What is the point of postulating phenomena you can never confirm (Cohen, Dennett, & Kanwisher, 2016a, p. 644)?

So why prefer my WHM? First, the illusion of richness account is based exclusively on the peripheral vision of Florence cases. I have found no studies that deal with ensemble statistics in foveal Ben cases. Crowding does not seem to occur in foveal vision (Alvarez, 2011, p. 127). While the illusion of richness account does not impoverish experience excessively, ensemble statistics do not at all describe Ben's foveal phenomenal FOC content, or his confident Seeming-that he sees every speckle richly. This failure is emphasised when we consider the difference between Ben's direct visual experience of a hen he is now looking at, with Ben's eyes-closed memory experience of a hen he stared at one minute ago. The latter case sounds very much like an image built from very sparse data and shares the characteristics of ensemble statistic perception. When Ben attends to the speckles in his memory-

²⁵² Whitney and Leib (2018, pp. 120–121) list the many general advantages of the concept of ensemble perception.

image, his confidence in seeing individual speckles—like Florence—drops.²⁵³ Remembered or imagined foveal images behave like Florence’s peripheral content,²⁵⁴ all the more starkly highlighting the contrast to direct foveal perception, and underlining the inadequacy of illusion of richness accounts to explain the latter.

Second, in 7.4.2.2.1 I argued that due to the epistemic obstacle to identifying phenomenal FOC that is not further processed, the evidence of Ben’s Seeming-that is all that we can reasonably hope for. Here I would add that the epistemic obstacle cuts both ways. It equally frustrates attempts to confirm that phenomenal FOC is in fact sparse, as the overflow sceptic would have it. If anything, the only report we have from Ben is his confident Seeming-that he sees every speckle richly. As I argued above, the burden of proof must surely be on those who question that. Nonetheless, this epistemic obstacle counts against both proponent and sceptic equally. While much more could be said, on balance, then, these considerations make the WHM preferable to the illusion of richness account in Ben cases.

8.6.2 Expanded Attention Account

The illusion of richness account circumvented overflow by rarefying conscious content to match the sparsity of attention. Another way around overflow is to expand the richness of attention to match the richness of consciousness—an *expanded attention* account. There are a number of ways of applying this idea. Here I consider three: redefining capacity so that Ben’s speckles fall within the capacity of WM; taking Ben’s Seeming-that as higher-order processing of every speckle; and

²⁵³ But not for the same reasons as Florence. Ben may, for example, be able to hold a clear picture of one or two speckles with eyes closed, but at the cost of the others. Florence is not capable of holding any of her peripheral trees clearly, even with her eyes open.

²⁵⁴ For a *predictive processing* account of the inchoateness of imagined content, see Clark (2016, Chapter 3).

redefining the boundaries of Executive Attention so that it is no longer limited by the capacity of WM. I consider each in turn but find them inferior to the WHM account.

8.6.2.1 *Capacity of Executive Attention in WM is Never Really Exceeded*

One approach is to measure capacity differently, thus bringing Ben's rich phenomenal content within the limitations of Executive Attention in WM. One might say that one hen is just one object—even with all its speckles—which in no way exceeds WM's capacity for object attention. This will only work, however, if one incorporates the fragile very short-term store into WM itself. One might say that rich content can be held in, say, the episodic buffer, or even the visuo-spatial sketchpad of BH, but only for a fraction of a second. This allows the Executive Attention of the central executive to operate upon it and make it phenomenal without operating outside WM as such.

But there are serious problems with this approach. First, while it's true that one hen is one object, a hundred *speckles* certainly exceed WM's capacity for object or feature attention. We know this, because Ben can manipulate the few speckles he does attend to in ways he cannot manipulate all the speckles at once (e.g., subitise, compare, etc.). Second, there are good empirical reasons for distinguishing between WM and iconic memory (8.4.2.2). To simply do away with this distinction for the purposes of solving the problem of overflow is somewhat *ad hoc*. It is true that on the Mushiness Principle (8.5) we should not be too dogmatic about borders between systems, and the borders in the WHM are very fuzzy, but this in no way discounts the fact that content in WM and iconic memory behave in very different ways. At any rate, so long as there is phenomenal content that is not processed in any sort of higher-order way, we have $C \sim A$, regardless of whether one locates that content within WM or without.

8.6.2.2 *Seeming-that is Executive Attention*

Another way for the sceptic to explain away overflow is to argue that the very act itself of generating Seeming-that about the phenomenal FOC renders it Executively Attended (and therefore not overflowing), since that is a kind of higher-order processing. Ben's phenomenal FOC is therefore *not* unprocessed in a higher-order way—it has been processed to produce the Seeming-that.

This approach is undermined by an application of the Target Question. The Targets of the FOC and the Seeming-that are different. The Seeming-that is merely a pointer or tag that lacks anything like the richness of the FOC itself. Thus, the Target of the Seeming-that is not identical to the Target of the FOC. For this sceptical argument to work, Ben's Experience-of every speckle is what would need to be identical to his FOC, which, we have seen, is simply not the case.

8.6.2.3 *Redefining Executive Attention*

The strongest option open to the sceptic here, I believe, is to challenge my definition of Executive Attention as being limited to just the executive of WM. The discussion of the capacity limitations of attention pointed out that unlike WM, in some circumstances (e.g., tracking multiple objects, 8.3.2) Executive Attention seems to be of virtually unlimited capacity. Perhaps, then, even if consciousness overflow WM, Executive Attention can in some way also overflow WM (as suggested in 7.4.2.2.3) to encompass all the rich content including, say, that in iconic memory. On this account Ben is confident in his Seeming-that he sees the speckles richly precisely because his Executive Attention is what brings all that rich content into consciousness. What limits his ability to report richly is not the capacity limitations of Executive Attention, but those of WM.

Again, the Mushiness Principle (8.5) dictates that we not be too dogmatic about drawing boundaries around systems like the cognitive executive. There are certainly precedents for holding Executive Attention to operate both within and beyond BH's central executive-episodic buffer axis (Baars & Franklin, 2003; Carruthers, 2017; Mogensen & Overgaard, 2018; Naccache, 2018). And Sergent and Rees (2007) argue that reporting is a signature of WM, but not necessarily of the global workspace. I invoked the executive of WM to provide a principled way of demarcating Executive Attention from Liberal Attention (on which Q is trivial), but there is no reason why the global workspace should not provide an equally principled way to establish this demarcation. Indeed, Block's espousal of overflow has been expressed explicitly in terms of the global workspace: he has argued explicitly that "the capacity of phenomenology is greater than the capacity of the global workspace" (Block, 2007, p. 549).²⁵⁵ I cautioned in 3.3.3 against begging the question of Q by simply defining attention via consciousness—in this case, defining the broadcasting of the global workspace as being what makes content conscious.

But the chief problem with this version of the expanded attention account is that it merely shifts the boundaries without in any substantive way accounting for the presence of FOC that is not further processed in a higher-order way. It is true that Executive Attention—however one demarcates it—is implemented in the production of many aspects of perceptual phenomenal content. It modulates visual content in accordance with our motivations, intentions, task requirements, etc. For example, spatial and feature attention interact to modulate the perceptual resolution of features in different locations (van Es, Theeuwes, Knapen, & Ré, 2018). Both exogenous and endogenous attention improve task performance both peripheral and central, albeit with complex patterns of modulation of visual resolution (Barbot & Carrasco, 2017). But these kinds of implementations are involved in producing higher-order content rather than FOC. That may be a large part of our ongoing stream of consciousness, but it only takes a small subset of that stream to be unprocessed in such ways to establish $C \sim A$. And that is exactly what most of Ben's

²⁵⁵ Fazekas and Nemeth (2018) and Pitts et al., (2018) relate these opposing views to neural theories of consciousness.

speckles are. He can see them richly without any overlay of attentional modulation with respect to the broader context and without modulation of their resolution. The Target of these attentional modulations is the hen, or perhaps a few speckles Ben attends to, but *never* every single speckle simultaneously.

It is also true that some attentional modulation must be involved in generating even simple FOC, such as “filling-in” of the blindspot in the visual field (Ramachandran & Hirstein, 1997, pp. 434–437). But this is ongoing, subconscious Liberal Attention, not Executive Attention. Ultimately, expanded attention accounts cannot account for phenomenal FOC that is not processed in any higher-order way without begging the question and adopting a Phenomenal Definition of attention.

The foregoing discussion suggests that regardless of how one draws the boundary between Liberal and Executive Attentions—whether that boundary is defined by the global workspace or by the more capacity limited executive of WM—the strong possibility of foveal overflowing phenomenal FOC remains. In fact, none of the expanded attention accounts seem capable of explaining that possibility, a possibility made actual in Ben’s overflowing rich foveal speckle experience. There remains one more interesting alternative to the WHM to consider.

8.6.3 Inchoateness Account

An issue has come up as part of other discussions that deserves, I believe, to be addressed directly, and provides a third way to account for Ben and Florence’s experience without including overflow. I shall call it *inchoateness*, by which I mean

phenomenal content that is less detailed, vivid, or clear.²⁵⁶ When Florence keeps her eyes still but shifts her attention to the peripheral visual field, she discovers that she cannot make out individual trees (much less leaves on the trees) and even her ability to identify things like colours and textures is significantly diminished. She recognises the diminution by contrast to her foveal content, which is vivid, clear, and detailed. I argued that Ben not experiencing this diminution in clarity within the hen was powerful grounds for thinking that his phenomenal content overflows his ability to Executively Attend to its detail. But might it be that Ben does indeed experience the speckles inchoately? If so, his foveal content may not exceed the capacity of Executive Attention, and his experience isn't substantively different to Florence's after all.

In this section I consider the two aspects that in general discourse define inchoate phenomenal experience and apply them to Florence and Ben. While inchoateness adequately captures Florence's peripheral forest experience, it fails to capture Ben's foveal speckle experience, and therefore fails to be preferable to the WHM overflow account.

8.6.3.1 *Two Aspects: Peripherality and Indistinctness*

The concept of inchoate phenomenal experience, as discussed in the literature, has two dimensions or aspects: peripherality and indistinctness. *Peripherality* recognises that all content is not attended equally. There is content that is the current focus of attention, and content that lies outside that focus. By and large, the focus is taken to gradually fade into peripherality rather than having a sharp boundary, much like the umbra and penumbra of the shadow cast in a lunar eclipse—hence, Koch and Tsuchiya's (2007) concept of the “near-absence of

²⁵⁶ I am not, here, adopting a question-begging Phenomenal Definition of attention, but describing a contingent feature of phenomenal content.

attention.” There are peripheralities other than that of attention—the central foveal and peripheral retinal fields, and the phenomenal umbra and penumbra (De Sousa, 2002; Galin, 1994; James, 1890; Mangan, 2007; May, 2004) being the most relevant two to this discussion.²⁵⁷

Indistinctness I take to be about the clarity of cognitive content.²⁵⁸ An example of indistinct perception is two-point tactile discrimination from the skin of the upper arm (6.3.5.1) where the indistinctness arises from very low-resolution sensory inputs due to the sparsity of skin receptors. Indistinctness may also arise from the content being processed relatively less—a lack of attention. This in turn may be because top-down attention is not directed to that content, or because of capacity limitations in Executive Attention and/or WM.²⁵⁹

Peripherality and indistinctness usually go together, as do centrality and distinctness, but they can dissociate. It is possible for content in the focus of attention to be indistinct, as when a subject tries very hard to discriminate whether one or two pins are touching her upper arm. She may vividly experience the inchoate character of the experience, but that is metacognitive—the tactile perception of her skin *itself* remains inchoate.

Similarly, peripheral content can be quite distinct. In vision, the foveal field is the focus of attention most of the time, resulting in the most distinct contents of our

²⁵⁷ Crick and Koch (2003, p. 124) suggest a *neural* penumbra to the neural correlates of consciousness that is not itself conscious, but may be responsible for “implicit priming.” The phrase “penumbra of consciousness” has popped up in other contexts (e.g., Lukens, 1896; Singleton, Mason, & Webber, 2004) where it seems to mean something more like implicit beliefs or attitudes that are not reflected upon by a subject.

²⁵⁸ As opposed to a diminution in the *phenomenality* as such of patches of the experience (recall the phenomenality-content distinction, 2.5). Gurwitsch’s field theories of attention and consciousness are of interest here (Arvidson, 2003; Yoshimi & Vinson, 2015), but this is a complex issue I bracket as not directly relevant to my argument.

²⁵⁹ Compare Watzl’s (2011a, pp. 151–153) discussion of *determinacy* and its relation to attention and consciousness. Most of his discussion relates to Florence cases, not Ben cases.

experience, while the retinal periphery is usually in the penumbra of attention, resulting in a rather indistinct peripheral vision. However, retinally peripheral content can become the focus of attention, and therefore become remarkably distinct. Astronomical observers learn how to keep the telescopic light of a “faint fuzzy” (an astronomical object such as a nebula or a galaxy whose faint luminosity is spread out and can only be detected by the more sensitive rods of the peripheral retina) fixated about 15° lateral of the fovea while shifting spatial attention to it. When foveated, the faint fuzzy is invisible. When attended in the rod-rich periphery, it magically becomes visible, often in prodigious (and picturesque) detail (R. N. Clark, 1990, pp. 17–18). This is an unusual case, and it takes time and practice for an astronomer to learn how to peripherally direct and maintain her spatial attention upon the retinal periphery so effectively as to examine a faint telescopic object in fine detail, and even sketch it with great fidelity.²⁶⁰ Two-point tactile discrimination and telescopic peripheral observation are clear evidence of the double dissociation between peripherality and distinctness.

8.6.3.2 *Back to Florence and Ben*

An inchoateness account of Florence’s experience is quite plausible, but can it be applied to Ben? First, consider Florence. While the umbras and penumbras of the attentional and *retinal* fields can dissociate, it is not so clear that those of the attentional and *phenomenal* fields can do so. Initially, when Florence’s retinal periphery is also in her attentional periphery, her peripheral phenomenal content is inchoate although it does not Seem-that it is. By holding her eyes still and shifting her focus of attention to the retinal periphery (like an astronomer) she revises her Seeming-that and concludes that her retinal peripheral content is indistinct after all. There are two reasons for this: the peripheral retina is relatively sparse in receptors; and her focus of attention brings this sparsity to her knowledge (Seeming-that). A

²⁶⁰ Bottom-up attention probably contributes—there is a “pop-out” effect involved.

third possible reason is not relevant here: the capacity limitations of Executive Attention in WM are not exceeded, because the content is too sparse to approach them anyway.

In Ben's case, we have the mirror opposite of Florence's case. The first two reasons cannot apply to Ben, while the third one might. For Ben, the foveal retina provides very rich inputs. And whether or not he is focusing his attention on his foveal field, it seems that he sees every speckle richly, a fact he can confirm by focusing his attention on the hen, but has no reason to doubt even if he shifts his attention to the retinal periphery. But here, the capacity limitations of Executive Attention in WM come into play.

There are two ways this aspect of the sceptical account can be formulated. First, even when Ben focuses his Executive Attention on the whole hen, he is not really phenomenally experiencing every speckle richly, but only inchoately, enough for phenomenal content not to overflow Executively Attended content (call this *foveally attended foveal inchoateness*, FAFI). Second, if Ben shifts his Executive Attention to the retinal periphery (like an astronomer), the sceptic would expect his foveal phenomenal content to become more inchoate, since it is Executive Attention that renders it distinct, and Executive Attentional resources have been largely shifted away from the foveal field (*peripherally attended foveal inchoateness*, PAFI).²⁶¹ I consider each in turn.

In the FAFI case, Ben's foveal phenomenal content is allegedly to some degree inchoate, much like the inchoate content of upper-arm two-point tactile discrimination, albeit for a different reason. Like tactile content, foveal content can be inchoate due to sparse inputs like Gabor patches on the edge of discernibility or the vague images of an imagined hen with one's eyes closed. But here the mechanism cannot be sparse input—it must be sparse processing power. The

²⁶¹ This would not be the case if Executive Attentional resources heavily exceeded what is required for making the speckles phenomenal. This seems unlikely, but I have no empirical evidence to back this up.

overflow proponent's response to FAFI is the one given in Chapter 7: the immediacy of direct foveal phenomenal FOC is a powerful datum that discounts the claim that the speckles are at all inchoate for Ben.

The PAFI case is more interesting. If it were the case that when Ben shifts his spatial Executive Attention to the periphery, the foveated speckles become inchoate, that would be a powerful piece of evidence for the overflow-sceptical case. It would show that phenomenal content is indeed only as rich as Executive Attention allows. When that Executive Attention is withdrawn (to the periphery in this case) phenomenal content is proportionately impoverished. But if it were found that shifting spatial attention has *no effect* on the phenomenal richness of the speckles, that would be powerful evidence for the overflow-proponent's account. Phenomenal richness does not vary with variations of Executive Attentional resources.

How might we go about determining whose account is right here? Foveal change blindness or inattentional blindness paradigms in the PAFI condition aren't going to help, because they measure *report* (and therefore Executive Attention) rather than the richness of phenomenal FOC. If Ben is asked to fixate his foveal field on the hen while shifting his spatial attention to a stimulus in the retinal periphery, and the hen is subtly changed while he does so, his inability to recognise the change is a failure of comparison—which requires Executive Attention—but not necessarily an indication of phenomenal sparsity.²⁶²

To provide an answer, we are going to need some kind of elaborate introspective report paradigm (Ramsøy & Overgaard, 2004). The simple dichotomic measures of Target perception typically employed will not suffice—detailed information about the phenomenal distinctness of the Target must be obtained (Haun et al., 2017, pp. 2–3). We want Ben to fixate his spatial attention on a suitably displaced retinally-peripheral distractor but at the same time introspect his foveal phenomenal content

²⁶² Watzl (2011a, p. 151) discusses some empirical work that shows that things like perceptual contrast improve with attention, but these involve higher-order processing performance rather than the phenomenal FOC that interests me here.

to produce a Seeming-that report about how distinct (and rich) it is. This is a complex procedure. Astronomers usually need to train to master the art of peripheral inspection of faint fuzzies. Subjects for this kind of experiment too would need to be trained. The possibility of inflation or filling in would need to be controlled for (Odegaard et al., 2018).

I have not found any empirical studies of this kind,²⁶³ and the epistemic woes above (7.4) cast some doubt over whether even the proposal just described would truly settle the matter. But on the basis of my arguments for the veridicality of immediacy of phenomenal FOC, it certainly seems worth trying. My own personal experiments suggest that no matter how strongly I focus my spatial attention to a peripheral Target, the *distinctness* of my foveated Target never changes. There is no proportional diminution in distinctness as a speckle becomes less and less attended. When Ben focuses his spatial attention peripherally, he is not counting the speckles, or noting their shape, or subitising a few of them, or comparing two of them to each other. Indeed, his metacognitive “awareness” of them may even be severely diminished. But so long as his lenses’ focal length doesn’t change, he is phenomenally *seeing them clearly*.

If this conclusion is borne out by empirical investigation, it would be a powerful reason to reject the inchoateness account of Ben’s experience: there are indeed Targets that are distinctly phenomenally experienced without being in any way Executively Attended.²⁶⁴ Again, this is an inference to the best explanation. The WHM account explains Ben’s continued foveal clarity while the inchoateness account cannot.

²⁶³ A lacuna confirmed by Anina Rich (personal communication). Weldon et al., (2016) demonstrates the effects of foveal content on peripheral discrimination tasks, and Eriksen and Eriksen (1974) flanker tasks and Lavie (2010) load manipulations demonstrate the effects of peripheral distractors on foveal *task performance*, but task performance measures Executive Attention, not phenomenal content. What we need here is to identify the effect of peripheral distractors on the character of foveal phenomenal perception.

²⁶⁴ Seeming-that excepted, see 8.6.2.2 above.

8.7 Chapter Summary

In this chapter, I posed the Capacity Question in relation to attention, consciousness, and WM. I found that they differ significantly, not only in their quantitative capacities, but also in the nature of their limitations. Whereas both WM and Executive Attention exhibit a trade-off between the number of items and the complexity of items, phenomenal consciousness does not. When these three capacity limitations are triangulated, we find a model of cognition emerging on which phenomenal content need not emerge solely from the implementation of Executive Attention in WM. This Witches' Hat Model, when compared to competing overflow-sceptical accounts (illusion of richness, expanded attention, and inchoateness accounts), better explains some crucial pieces of evidence, especially Ben's Seeming—that he phenomenally experiences every speckle on the hen richly and distinctly, regardless of where his attention is focused, and despite only being able to higher-order process and report a small subset of that rich content.

I am not here arguing that because the WHM is correct, consciousness overflows attention. Nor am I arguing that reverse, that because cases like Ben's hen show that consciousness overflows attention, therefore something like the WHM must be correct. I am not convinced that I have enough evidence to make either of those cases. What I am confident in arguing is that cases like Ben's hen are not only plausibly explained by conscious overflow on the WHM, but that my explanation has many reasons to recommend it over competing explanations, at least at this stage of our knowledge.

In the final chapter, I consider how the results of this quest for dissociated attention and consciousness in Part II can be situated within my Set Theoretical Framework, and therefore, what they mean for Q.

9 **Answer(s) to Q**

9.1 Conclusions, Implications, Applications.

Is attention both necessary and sufficient for consciousness? I have argued in this treatise that this question, Q, is not one question, but many. It can be addressed analytically or empirically. It depends on what exactly one means by attention and consciousness. Each way of posing Q admits of its own answer, and reveals different things about the relationship between the two. This treatise therefore focused on a quest to answer one of the most interesting versions of Q: *is Executive Attention both necessary and sufficient for phenomenal consciousness?*

In this final chapter, I bring together the conclusions of the preceding chapters, consider the implications for the discourse surrounding Q, and conclude by sketching some broader implications and applications of my work.

9.2 Seeking Answers

In this section, I first summarise the foundational basis I developed in Part I: definitions of attention, consciousness, and WM; analytical and empirical approaches to Q; the Set Theoretic Framework (STF) that describes Q's conceptual space; the kinds of relationship that might underlie the pattern of co-occurrence; and the spectrum of weak to strong readings to which an answer to Q may pertain. For the purposes of comparison, I then pose three different versions of Q and discuss them in the light of the aforementioned foundations, before returning to draw detailed conclusions about the version of Q (Executive Attention and phenomenal consciousness) that has been the chief topic of Part II. I conclude this section with some observations about the third concept in the triad—WM.

9.2.1 The Terrain

We can think of addressing Q as being something like adjusting the dials on a panel and reading off the resultant outputs on a series of displays. Each dial represents a particular variable—definitions, analytical or empirical approach—and each output displays represents a value for a particular parameter—answer to Q, Scenario, kind of relationship, and readings. The terrain thus mapped out is truly vast, and I believe that virtually all the analytic or empirical literature on the relationship between attention and consciousness can be located somewhere on it. In this relatively brief treatise, I have of course only been able to explore a tiny portion of it, and there remains much beyond this portion that is either little explored or not at all.

9.2.1.1 *Definitional Variety*

One set of dials adjusts the different ways of defining attention and consciousness. In Chapters 2 and 3 I considered the many different ways the terms attention and consciousness have been used in the literature. One way to arrive at different answers to Q is to employ different definitions of attention and of consciousness and to determine the relationship between each pairing.

On the attention “dials,” one may select any of the kinds of attention: internal versus external; endogenous versus exogenous; voluntary versus involuntary; Liberal versus Executive; etc., and with any of the definitions of attention I catalogued in my Taxonomy: Behaviourist; Phenomenal; or the many kinds of Mechanistic definitions. On the consciousness “dials,” one may select any of the different definitions of consciousness described in 2.3: physiological; access; phenomenal; etc. Each of this large array of settings will produce its own outputs. On our current state of knowledge, some will produce more confident outputs than others.

9.2.1.2 *Analytical and Empirical*

In 1.4.1 I described analytical and empirical versions of Q. This treatise leads us to the conclusion that while on most definitions of attention and consciousness, the two concepts can certainly come apart analytically, an answer derived from the empirical evidence is less clear. I employed both analytical and empirical approaches in this treatise.

9.2.1.3 *Patterns of Co-occurrence: STF*

The simple yes or no answer to Q may be further elaborated by identifying the pattern of co-occurrence according to my STF. In 4.2.1 I described the four possible

Combinations of attention and consciousness in any given cognition: $A \& C$; $A \sim C$; $C \sim A$; and $\sim A \sim C$. These Combinations allow for local answers to Q—whether attention is necessary and sufficient for consciousness in a particular cognition. Of chief interest, however, has been the answer to Q that characterises a whole cognitive economy. For this, we need Scenarios. To that end, I developed an exhaustive framework for all the possible Scenarios (sixteen) that might relate attention to consciousness in a cognitive economy. Of these, only four were plausible or likely to obtain, although the other twelve are not without interest. My literature review in the same chapter revealed that each of the four has its proponents.

9.2.1.4 *Kinds of Relationship*

The yes or no answer to Q and the Combination or Scenario may be further elaborated by elucidating the kind of relationship that underlies the pattern of co-occurrence of attention and consciousness, as I sketched out in 4.4. The possible outputs here are the Identity Hypothesis (IH), Partial Constitution (PC), Causation (CA), and Mere Correlation (MC), refined by their many sub-varieties. This question opens the door into more far-reaching questions about the very natures of attention and consciousness.

9.2.1.5 *Weak and Strong Readings*

Finally, the significance of the output—or perhaps the domain over which it obtains—can be described by which of four Readings applies to it (4.2.3.2): Very Weak; Weak; Strong; or Very Strong. The preponderance of material covered in this treatise relates to human cognition, often normal, but sometimes pathological. My answers to Q and preferred Scenarios below should therefore be taken as a Weak

Reading—applying to human cognition generally, though with some possible unusual exceptions.

9.2.2 Different Questions, Different Answers

Having mapped out the terrain of possible inputs and outputs in addressing Q, I turn now to some of the more interesting possible answers that might be derived by twiddling the dials.

9.2.2.1 *Phenomenal Attention and Phenomenal Consciousness*

In 3.3.3, I coined a class of definitions of attention that I called Phenomenal Definitions of Attention. These were definitions that included phenomenal consciousness as an intrinsic feature of attention. Clearly, on Phenomenal Attention, there can be no $A \sim C$, since any cognition that lacks phenomenality would thereby lack an integral constituent of attention, and cannot be attention. As I have often remarked, this trivialises Q by begging the question, by defining consciousness into attention *a priori*.

Phenomenal Attention eliminates the possibility of $A \sim C$, but it leaves open the possibility of $C \sim A$. For example, if attention is considered to be the “chief tenancy of consciousness” (as per Bradley, 3.3.3.1) one may still ask whether there are lesser tenants. It may be that there is conscious content altogether peripheral to the focus of attention and therefore altogether unattended, which is local $C \sim A$. Or it may be that attention gradually fades in magnitude as one moves outward (not necessarily spatially) from the focus of attention, so that all conscious content is attended, but peripheral content is only minutely attended, KT’s “near-absence of attention”

(4.2.4.2). On this view, there is no true $C \sim A$. The relationship between Phenomenal Attention and phenomenal consciousness admits two possible Scenarios, then, depending on which of the two foregoing views one prefers: $C \supset A$ (no to Q—attention is sufficient but not necessary for consciousness) or $A = C$ (yes to Q), respectively.

9.2.2.2 *Liberal Attention and Phenomenal Consciousness*

I also coined another way of defining attention, Liberal Attention, on which any implementation of any attentive strategy (as catalogued in the Operations Definitions of attention) is enough to constitute attention. Whereas Phenomenal Attention begs the question by making consciousness intrinsic to attention, Liberal Attention begs the question by making attention ubiquitous, such that virtually any cognitive process at all (and many non-cognitive processes as well) involves attention. For Liberal Attention and phenomenal consciousness, there can therefore never be $C \sim A$.²⁶⁵ But there will almost certainly be $A \sim C$, since it is widely accepted that much of the brain's work goes on unconsciously. Adopting a Liberal Definition of attention commits one to the Scenario $A \supset C$ (no to Q—attention is necessary but not sufficient).

9.2.2.3 *Executive Attention and Access Consciousness*

Both Phenomenal and Liberal Attention are not as interesting as the varieties of Q addressed by most of the literature—those involving what I have called Executive Attention. But the variety of consciousness involved is not always so easy to determine. Block's valuable (though somewhat controversial) categories of access

²⁶⁵ Unless one accepts the possibility of consciousness independent of cognitive processes.

and phenomenal consciousness cover many of the discussions in the literature. But these two kinds of consciousness produce different answers to Q.

I argued in 7.2.2 that access consciousness is roughly equivalent to Executive Attention. Defining consciousness in Q as access consciousness therefore begs the question in a way not unlike Phenomenal Attention: rather than defining attention as a kind of consciousness, it defines consciousness as a kind of attention, at least on Block's definition of access consciousness and my definition of Executive Attention. On this version of Q, then, there can be no $A \sim C$ or $C \sim A$ since Executive Attention roughly just *is* access consciousness, and the only viable Scenario will be $A = C$ (yes to Q). If one considers the concepts of access consciousness and Executive Attention not to be equivalent, then there will be the possibility of $A \sim C$ and/or $C \sim A$ depending on how they diverge, thus opening the door to other Scenarios.

9.2.2.4 *Executive Attention and Phenomenal Consciousness*

To my mind, the *really* interesting version of Q—and the one with most promise to illuminate the fascinating questions about the nature of cognition generally and attention and consciousness specifically—is the relationship between Executive Attention and phenomenal consciousness. It is to this version of Q that Part II has been devoted. In this section I summarise the arguments presented there and draw some final conclusions from them.

In Chapter 6.2 I explored the evidence for $A \sim C$, and found it quite compelling, a minority of dissenters notwithstanding. To the sample of existing empirical evidence, I added some other plausible candidates for $A \sim C$ that I have not found so far discussed in relation to Q in the existing literature. Further, those dissenting voices present ambiguous arguments against $A \sim C$, such as Prinz (6.2.2), whose acceptance of “unconscious perception” basically amounts to an acceptance of $A \sim C$.

This affirmation of $A \sim C$ provides a brute answer of “no” to Q. But the right Scenario has not yet been established. Given what I have taken to be the uncontroversial affirmation of $A \& C$ and $\sim A \sim C$, the choice of Scenario rests upon whether or not $C \sim A$ is instantiated in a human cognitive economy. The two remaining viable Scenarios are $A \supset C$ (if $C \sim A$ is not instantiated) or $A \cup C$ (if $C \sim A$ is instantiated).²⁶⁶ Keep in mind that we need only one undeniable case of $C \sim A$ to confirm a Scenario. That being the case, I embarked upon a thorough and methodical exploration of both the analytical congruence and the empirical evidence for $C \sim A$. Analytically, the foundational concepts described in Part I played an important role, especially my definitions of phenomenal consciousness as a what-it-is-like, situated, temporal, first-person perspective and of Executive Attention as a suite of strategies for structuring cognition implemented by the cognitive executive and sharing the characteristic of structuring content for further processing. This version of Q, then comes down roughly to this: can there be cognitive content that is phenomenal yet not structured by Executive Attentional strategies and further processed?

In Chapter 6, I then proceeded to critique five different possible classes of $C \sim A$ that have been discussed in the literature—Pure Consciousness, Global Unprocessed Content, Simplicity, Chaos, and Timing—and found that while many of them are suggestive, none of them provides conclusive evidence for $C \sim A$.

In Chapter 7 I turned to what I consider to be the most promising putative case of $C \sim A$, a particular kind of local Unprocessed phenomenal Content: phenomenal overflow. Here I drew the distinction between three kinds of content: *First Order Content* (FOC); *Experience-of*; and *Seeming-that*. I noted that FOC does not beg the question by being defined as phenomenal content unstructured by Executive Attention, since its conceptual possibility in no way guarantees its empirically confirmed existence. By way of analogy, to posit the idea of multidimensional strings is in no way to beg the question of whether String Theory is true.

²⁶⁶ These two Scenarios answer “no” to Q in different ways: $A \supset C$ means that attention is necessary but not sufficient for consciousness and $A \cup C$ means that attention is neither necessary nor sufficient for consciousness.

The chief competitor to the overflow interpretation of the apparent richness of experience is the *illusion of richness* account, on which the brain merely fills in a kind of averaged place-marker in the periphery of attention, underwritten by ensemble statistics. Florence thinks she sees every tree and leaf in the forest richly, but in fact, she sees only the content of a very narrow focus of attention richly, with the rest of her forest experience being a Seeming-that it is rich, without actual FOC or Experience-of rich detail peripheral to the focus of attention. I argued that this explanation finds support in the neurophysiology of the human visual system, from the retina through to the higher visual centres. However, Ben's speckled hen experience cannot be thus explained away, since it is the phenomenal content that lies within his foveal visual field and spatial focus of attention that is far too rich to all be processed by the executive.

There remains an intransigent epistemic obstacle to confirming this assertion. Our chief source of information as to whether content is phenomenal or not is report by the subject of the experience. Yet report necessarily requires Executive Attentional processing. If there is phenomenal content that is not reported (and unattended) then we can only know this through some method of identifying it without report. No-report paradigms have not yet been able to deliver this, and in fact, may well never be able to do so, since any no-report paradigm must initially be calibrated using subjective report (7.4.3.3).

Nonetheless, I mounted an abductive argument for overflow as the best explanation of the data available in two ways. First, there is the immediacy of phenomenal FOC (7.4.2). Ben's Seeming-that he sees every speckle on the hen is clear, confident, and only strengthened by further introspection in a way that Florence's Seeming-that she experiences every tree and leaf is not, and constitutes powerful evidence that he does indeed have a phenomenal FOC experience of every speckle, richly. Further, Ben's inability to accurately perform Executive Attentional process such as count or even subitise the dozens of speckles is not evidence that he does not experience every speckle. The illusion of richness account may explain Florence's forest (although, see 9.3.1.2), but cannot explain Ben's hen.

Second, in Chapter 8, I developed an argument from the capacity limitations of attention, consciousness and WM. WM capacity is limited by the cognitive executive's limited capacity to process content at any given time, the dimensions of which are reasonably well described (though not conclusively) in the empirical literature. Executive Attentional strategies implemented by the cognitive executive are therefore defined by this limited capacity of the executive, even if attentional strategies *per se* are not so limited. There is no way to enlarge the capacity of Executive Attentional processes, therefore, without enlarging the capacity of the WM executive, at least on my definitions of these concepts. On the other hand, phenomenal consciousness has, in theory, and like attentional strategies *per se*, no capacity limitations. A subject can in theory synchronically experience any amount of conscious content that are produced in the right way. What is more, both attention and WM exhibit a trade-off between the number of items and the complexity of those items, while consciousness does not. Even if one rejects my characterisation of Executive Attention as the attentional strategies implemented by the executive of WM and allows Executive Attention to operate outside of WM (in a Baars global workspace, for example), this difference in trade-off patterns remains. It would thus be remarkable indeed if the content of Executive Attention and phenomenal consciousness were always identical.

The question then becomes whether we think conscious content can be generated outside the executive of WM. I further argued that when we consider the structure of memory storage and manipulation in human WM, the picture that emerges is not of any discrete module of WM, but of an organic, graded interplay of stores where duration of storage varies inversely with storage capacity—my Witches' Hat Model (WHM, 8.5). If this is indeed the structure of our storage and manipulation systems, there is no principled reason to think that phenomenality is sharply restricted to content in the executive of WM, unless we think that it is the executive of WM itself that is solely responsible for making content phenomenal. But that would be begging the question. A final answer to the question of whether content in a rich short-term store can be definitively shown to be phenomenal awaits a convincing way for identifying specific conscious content in the brain. But I venture to suggest that my

account of phenomenal overflow within a WHM paradigm provides a more plausible account than the alternatives (illusion of richness, expanded attention, inchoateness) for Ben's Seeming-that he sees every speckle in the absence of an Experience-of every speckle.

9.2.2.5 *Summary of Conclusions for Executive Attention and Phenomenal Consciousness*

In 9.2.1 I identified a number of output parameters or dimensions on which we can answer Q. Here I summarise the answers that fall out of my analysis in the case of Executive Attention and phenomenal consciousness in 9.2.2.4.

9.2.2.5.1 *STF*

If phenomenal overflow does indeed constitute a *bona fide* case of $C \sim A$, then all four Combinations are instantiated in a human cognitive economy. The Scenario that best describes this cognitive economy is thus $A \cup C$ (no to Q—attention is neither necessary nor sufficient for consciousness).

9.2.2.5.2 *Kinds of Relationship*

My analysis does not point to a single answer to the question of the kind of relationship that obtains between Executive Attention and phenomenal consciousness, but it can rule some of the options out. Most obviously, the Scenario $A \cup C$ rules out the Identity Hypothesis. If attention and consciousness can be instantiated each without the other, then it is not possible for them to be identical. One might propose that they are the same entity that is called different things in different circumstances, but my careful and distinct definitions of both entities do not leave that door open. We can also rule out varieties of Partial Constitution on

which just one of the two partially constitutes the other, since both the Combinations $A \sim C$ and $C \sim A$ obtain. We cannot, however, rule out the variety of Partial Constitution where attention and consciousness each partially constitute the other. Mere Correlation is unlikely to be the answer, because of the intimate pattern of co-occurrence, which would be most odd if there were no constitutive or causal relationship of any kind tying the two together.

This leaves the possibility of some kind of Causation as being the most plausible. There seems little doubt that the same FOC can be Executively Attended or can become phenomenal. There seems little doubt that it is often (but not always) the case that attending to a target brings it into consciousness, or being conscious of a target facilitates attention to it. There are likely to be close connections between the neural correlates of attention and of consciousness. A very plausible case can be made for a *complex causal chain*, a variety of CA (8.4.1). But ultimately, my approach here was not designed to provide a clear answer to the question of the nature of the relationship, but only to Q. I cannot, on the material covered in this treatise, draw confident conclusions about the kind of relationship, any more than to sketch the constraints above.

9.2.2.5.3 Readings

The evidence discussed in this treatise related almost exclusively to human cognition, whether “normal”²⁶⁷ or pathological. The Scenario of $A \cup C$ ought therefore to be taken as obtaining on a Weak Reading—it is the Scenario that obtains in all human cases except some rare and extreme conditions. It is not a Very Weak Reading, since it seems to reflect the basic character of human cognition, and does not vary from individual to individual, except for rare, extreme cases. It is not a Strong Reading, since there are extreme exceptions such as vegetative patients in

²⁶⁷ Whether there is such a thing as “normal” cognition, and how exactly we might go about delineating its boundaries, are interesting questions (Holmes & Patrick, 2018).

whom there may be an absence of either Executive Attention or phenomenal consciousness altogether. And of course, it is not a Very Strong Reading since different Scenarios might obtain in the physiologies of different species, or in hypothetical alien species, or even in advanced artificial intelligence computers.

9.3 Working Memory

What role has WM played in my analysis? The literature on Q often employs empirical work on WM, suggesting that this body of research would be a rich vein to mine, as has proved to be the case. By adding WM as a third concept, triangulated with attention and consciousness, I was able to define the most interesting form of attention in relation to Q—Executive Attention—and distinguish it from Liberal Attention in a principled way that reflects both the nature of human cognition and much of the literature on Q, even if much of that literature does not itself accurately define what is meant by “attention.” WM and its relation to other kinds of memory, e.g., iconic memory, also provided a basis for my WHM of cognition which in turn provided the framework for my comparison of capacity limitations of attention, consciousness, and WM. WM models also underwrite my distinction between FOC, Experience-of, and Seeming-that. The latter two are the product of Executive Attention in WM, while the former is not (or at least, is not necessarily). My analysis of phenomenal overflow in Chapter 7 depended heavily on this distinction.

9.3.1 Models of Working Memory

My analysis also has important implications for models of WM, such as BH and Cowan's. on both of these influential models, content is assumed to become phenomenal only when it is attended by the central executive or the focus of attention respectively. If phenomenal overflow is indeed possible as I have argued, and if something like my WHM better captures what is happening, then these models of WM would need to be revised, at least in respect of the relationship between attention and consciousness.

For Prinz, attention is the gateway to WM. By this, he is thinking of attention as something broader than BH's central executive, which interacts with the slave components (visuo-spatial sketchpad, phonological loop, episodic buffer). But for Prinz, it is attention that is responsible for bringing content into the slave components of WM in the first place. This is a significant difference between Prinz's view of how attention works in WM and that of BH. But it is also incompatible with my view, in that it limits conscious content to only that which is attended (recall from 4.3.1 that I classified Prinz as subscribing to the $A = C$ Scenario). The possibility of phenomenal FOC in something like iconic memory is incompatible with Prinz's view that only the further processing of attention can make content conscious.

9.3.2 Other Pertinent Questions

In Chapter 5 I articulated a number of pertinent questions in relation to WM and its interaction with the rest of cognition, all of which I touched upon, but only one of which I was able to explore in significant detail—the Capacity Question in Chapter 8. The other questions all hold great promise for further investigation and future research. For example, the *Duplication Question* is a fascinating one that holds much promise for progress on Q. While we have some idea as to how a memory is encoded, maintained, and retrieved, we still know agonisingly little about the neural

substrate of WM. The unresolved conflict between BH's components and Cowan's activated long-term memory models reflects this ignorance. For example, should it be shown that content only becomes phenomenal when a duplicate trace of that content is encoded into some WM neural buffer, that would be a significant blow to my account, making the possibility of phenomenal FOC in iconic memory less appealing. If on the other hand it were to be shown that something like neural reverberation operating at the neural locus of iconic memory stores is the explanation for how such content becomes phenomenal, that would strongly support my account. Below, I reflect briefly on the *Metaphysical Question* (9.5.1) and the *Function Question* (9.5.2.2).

9.4 Implications for the Current Discourse on Attention and Consciousness

In this section I highlight some further ways—beyond the actual conclusions I drew above—in which my analysis might impact on the discourse over the relationship between attention and consciousness. I consider some further ways that the foundations I developed in this treatise might be of use; I explore some further ways that the search for $C \sim A$ might be executed; and I draw a somewhat startling consequence from my account—the nature of our phenomenal inner lives may be quite different to how we commonly conceive of them.

9.4.1 Foundations

In Chapter 4, I developed a number of conceptual tools to facilitate my investigation of Q. These included my Set Theoretical Framework (STF) with its exhaustive possibilities for Combinations and Scenarios, the catalogue of four possible kinds of relationship, and the different Readings of an answer to Q. Further, in 1.4.6 I proposed four salient questions—Target, Timing, Variety, and Consequences—to assist in assessing whether any given empirical paradigm and its experimental results actually addresses Q or not. These are general tools—whether in their current form, or in improved forms—that may be of assistance in bringing clarity to the discourse surrounding Q, and perhaps even stimulate new approaches to the question. They may also be useful for questions other than Q.²⁶⁸

I observed earlier that definitional ambiguity has been a major bane on this discourse. To that end, I venture to hope that my approach to defining phenomenal consciousness in terms of its core characteristics, and especially my definitional taxonomy of attention, would be of some use in circumventing that ambiguity. Both of these are, of course, also subject to improvements. I have not insisted on any one way of defining attention (although I have adopted one—Executive Attention defined operationally—for the discussion in Part II), but have rather provided an integrated range of options. By being clearer on what exactly any given account means by the terms attention and consciousness, and by situating it against my STF, kinds of relationship, and Readings, it is to be hoped that more fruitful comparison, engagement, and interaction may be facilitated between different accounts.

²⁶⁸ For example, the STF approach can also be extended to more than two concepts. Applied to *three* concepts (e.g., attention, consciousness and working memory), there are eight possible Combinations and 256 possible Scenarios. More generally, the number of Combinations possible with x concepts is 2^x , and the number of Scenarios possible with n Combinations is 2^n . Such extensions quickly become unwieldy, but may still be useful in some ways.

9.4.2 Consciousness Without Attention

In 6.2.2 I only briefly surveyed candidates for $A \sim C$ and suggested some new ones. But the survey of possible candidates for $C \sim A$ that followed is, I believe, the most comprehensive to date, although the broadness of its scope meant a certain lack of depth in the treatment of individual candidates. No doubt my conclusions on some of the candidates are open to challenge, but I hope that so broad a survey bringing together so many candidates in one place may be of use to other researchers as a reference, and my classes of candidates for $C \sim A$ may perhaps also be an inspiration for new places to look for candidates within those classes, or even for new classes altogether. I also identified some candidates (e.g., focal epileptic seizures, 6.3.4.2) where a little empirical investigation of the right type may yield significant progress.

The one candidate I did explore in depth—phenomenal overflow—has been a hot topic in recent years. My WHM is not significantly original in its concept, but is a new way of expressing the idea that content can become phenomenal outside of Executive Attention. The example of Ben's foveal hen highlights a case that has hitherto received very little attention in the literature (especially empirical), despite it being, I believe, a stronger candidate for overflow than the more common Florence's forest case. More empirical work on overflow in Ben cases would be especially interesting, although the many methodological obstacles inherent in investigating it means that new and creative approaches are needed to probe it.

My triangulation of capacity limitations of attention, consciousness, and WM (Chapter 8) provided the basis for an account of phenomenal overflow that was more plausible (I hold) than its sceptical rivals. Consideration of capacity limitations can profitably be applied to many other candidates for $C \sim A$, another avenue for future research. What is more, the triangulation approach itself can be profitably implemented to address Q using other cognitive constructs in relation to attention and consciousness. How does *long-term memory* relate to Q? What of *cognitive control*, or *motor command*?

The kind of thorough analysis I employed in Chapters 7 and 8 might also be applied to the candidates for $C\sim A$ in Chapter 6. Some of the conclusions I reached in that chapter might thereby be revised. Phenomenal overflow is a local version of the class of “unprocessed content” (6.3.4), but some of the arguments applied might be adapted to both local and global members of this class.

Throughout Part II I have contrasted Florence and Ben cases to strengthen the case for overflow in the latter. However, a more thorough analysis of Florence cases might also lead to a revision of the conclusions I reached about them. I doubt that an illusion of richness account for Florence cases could be overturned completely, but there may still be room for finding some form or degree of overflow there. For example, the foveal field is quite small in relation to the whole visual field (Appendix 7), certainly smaller than the subjective field of focus it Seems-that we experience. That is, our Seeming-that we have not only a rich foveal experience, but an equally rich close-parafoveal experience is nearly as confident as Ben’s confidence in his foveal richness. Certainly, saccades and other factors may be contributing to that richness, but there is scope here for something like the immediacy of foveal richness, albeit attenuated as we move further away from the fovea. There are interesting parallels between Ben’s rich experience of speckles and Florence drinking in the broad vista of a dark night sky awash with stars. Exciting developments in experimental paradigms may soon lead to more confident answers in these cases (e.g., Kentridge, 2011, pp. 233–235; F. F. Li, Vanrullen, Koch, & Perona, 2002; Vandenbroucke et al., 2014), so any advance on the analytical front would be advantageous.

9.4.3 Forgotten Conscious Content

Finally, my account of phenomenal overflow, if correct, would support what is currently a minority interpretation of some prominent empirical paradigms. On my account, it is quite plausible that subjects who experience inattentional blindness,

change blindness, or flash suppression do indeed have a phenomenal FOC experience of the content that is masked, but they simply cannot report (in the broad sense above) that content (Wolfe, 1999a).²⁶⁹

Can we say that Ben's hen experience differs from his hen-minus-one-speckle experience (change blindness)? Yes and no. Yes, because his phenomenal FOC experiences were different. No, because Executive Attention was not able to *compare* and tell the difference (Lamme, 2010, pp. 219–220; Pinto et al., 2017, pp. 223–224) so that the two Experiences-of and Seemings-that were identical. Like the less interesting passages in the novel, the fleeting experience in iconic memory is soon forgotten and lost forever. This may be analogous to “white dreams” where, upon being awoken from a dream, a sleeper reports a confident Seeming-that she had dreamt richly but is unable to recall any content of the dream (FOC/Experience-of) whatsoever (Siclari et al., 2013; Windt et al., 2016, p. 879).²⁷⁰ In cases of severe memory pathology—such as the memory loss of Alzheimer's disease (Förstl & Kurz, 1999), or the classic extreme case of HM who was unable to lay down new memories (Corkin, 2002)—not even a Seeming-that persists to mark the lost phenomenal content. Yet in all these cases we have no reason to doubt that the initial phenomenal content was in fact experienced by the subject.

²⁶⁹ An idea foreshadowed by Condillac: “Those perceptions that we attend to can seem to drown out the others and produce the illusion that they alone exist, whereas those perceptions that we are less conscious of can be so faint that it is impossible to recall that we have had them the instant after the stimulus that produced them fades” (Falkenstein & Grandi, 2017, sec. 5). For modern evidence of rapid overwriting of phenomenal content, see Landman et al., (2003, p. 150). For recent evidence and arguments against this view, see Ward (2018), Pitts et al., (2018, pp. 2–3), and Matthews et al., (2018). Ramsøy and Overgaard (2004, p. 20 Note 1) observe that while in the cognitive science literature, subliminal perception is held to be completely without phenomenality, in the phenomenological literature, some kind of minimal phenomenality is associated with subliminal perception. Space prevents me from engaging in a detailed analysis here.

²⁷⁰ For an alternative “illusion of richness” type interpretation of white dreams, see Fazekas et al., (2019).

I have not mounted a case that this *must* be what happens in masking conditions, only that it is a quite plausible possibility. Ultimately, my arguments are arguments for epistemic humility, for admitting we cannot yet be sure of answers, and should keep an open mind. If this kind of fleeting consciousness is what happens, however, there are serious implications for our understanding of conscious experience. The idea that my conscious experience might include so much content that is so fleeting as to be immediately forgotten as though it never happened threatens our sense of the nature of our selves. To realise that you cannot trust the content of your own *senses*—as illusion of richness accounts would have it—is disturbing enough. But to realise that you do not know what you are internally *experiencing* is utterly disconcerting. It sounds like a paradox from Plato. How can I not know what I know? There is much fertile ground for philosophical exploration here.

9.5 Broader Implications and Applications

The implications of my account extend beyond Q. The issues I tackle below are treated only very briefly, and only insofar as they touch upon my analysis of Q, although they are, of course, far broader and deeper than I can even hint here. In this section I consider four broad areas of implication and application: the ontology of attention and consciousness; theories of consciousness; other cognitive functions; and some ethical applications.

9.5.1 Ontology

The question of what kind of metaphysical entities attention and consciousness are has received many answers over time. Here I consider some possible light that my account might shed on it, and upon how attention and consciousness fit into the rest of reality.

9.5.1.1 *Attention*

My quest to distil the essence of attention in Chapter 3 led me to classify attention as a strategy, which would fall under the metaphysical category of an abstract particular (Lowe, 2002, p. 16). But it is a particularly interesting particular. On my definition, there is no reason to think that attention is limited to human cognition, or to cognition generally, or even to biology. A plant slowly turning its leaves to face the sunlight fully implements the attentional strategies of orienting, tracking, selection-for-action, and control. So does a security camera fitted with servo motors that allow it to pan so as to maintain a moving object in the centre of its optical field. Liberal Attention is not only virtually ubiquitous in human cognition, it seems, but extends far, far beyond it.

What does this promiscuity of attention mean? We could take this as a blurring of the metaphysical lines between cognition and non-cognition—between mind and non-mind. Many assume a sharp line between the two, but perhaps they melt into one another? We may even begin to wonder if the distinction is a valid one at all. This should not disturb us. It is often forgotten that brains are organic, and that they operate “mushily” (8.5). They do not have discrete parts performing discrete clockwork operations so much as vastly interconnected parts operating in chaotic, probabilistic ways. The operation of a brain is in many ways far more akin to the operation of a liver than it is to that of a silicon chip. The promiscuity of attention is plausibly a reflection of this organic basis of cognition, and we ought not be

surprised that other organic entities—e.g., livers and sun-following plants—also implement Liberal Attention.

What of non-organic attention? Can we take attention's promiscuity as evidence for reductionism about cognition? And if attention in a tracking security camera can be completely reduced to its parts, operations, and organisation (Bechtel, 2008), why not human attention, including Executive Attention? I think that a good case can be made for this kind of reduction in the case of attention, and even Executive Attention, and the case falls out naturally from the Mechanistic definitions of attention in my Taxonomy. Much more problematic, I think, would be to infer the reducibility of phenomenal consciousness from the reducibility of attention. If $A \cup C$ is the correct Scenario in humans, and $C \sim A$ obtains, then, as I argue below (9.5.2.3), attention cannot simply be substituted for consciousness to draw conclusions about consciousness.

9.5.1.2 *Consciousness*

Attention and consciousness, while capable of being defined as distinct entities, certainly seem to be intimately connected. Yet, there is an asymmetry evident throughout Part II. It was not difficult to establish $A \sim C$, even when the attention was Executive Attention. But why was it so difficult to establish $C \sim A$? I highlighted some prominent epistemic or methodological challenges in 7.4, but might there be something more to the matter than mere methodology? Might the connection be a causal one (CA)? What might this discussion tell us about the very nature of phenomenal consciousness? In this section I speculate rather freely on some possible directions this topic might take in light of my arguments, along the lines that attentional strategies are inherent in the cognitive "system" while phenomenal consciousness can plausibly be seen to be product of such systems, yet standing outside the system as such.

Mechanistic Attention is integrally involved in the processes that produce either the content of consciousness, or phenomenality itself. This means there is a significant ontological difference between attention and consciousness: being a strategy, attention is *intrinsic* to the system, in a way that consciousness is not, for consciousness is an *output* of the system. Attention *characterises* the mechanism of cognition (in the Bechtel sense). Executive Attention characterises the mechanism of the cognitive executive. But phenomenally conscious content is the result, consequence, or *output* of that system's operations. On my account, it is an output not only of the cognitive executive, but also of other, non-executive systems. Even if Liberal Attention is always involved in producing the *content* of consciousness, *phenomenality* itself may not require attention—e.g., on protopanpsychist accounts of consciousness.

This difference helps explain why consciousness is so much harder to *classify* ontologically than attention. Mechanisms, and strategies that characterise their operations, are relatively common in our experience of the world, and easily classified by analogy with other examples. But outputs come in many metaphysical kinds, and the output that is consciousness has few counterparts in nature or possibly none, as I hinted in Chapter 2.

The ontological difference between attention and consciousness is further highlighted by the contrast between the divisibility of attention and the unity of consciousness. In 8.3.2 I discussed the surprising potential for Executive Attention to be divided among multiple targets simultaneously. This now appears to be well established empirically. There are strong (though not conclusive) arguments, on the other hand, to think that phenomenal consciousness is intrinsically unitary—the Phenomenal Unity Thesis (Bayne, 2010; Bayne & Chalmers, 2003). Phenomenal experience is always a single unitary experience with multiple contents. What is more, the contents of a single experience can alter each other—the phenomenon of phenomenal interdependence (Dainton, 2000; Kaldas, 2015). But attention seems to be capable of having no such unifying feature when it is divided among multiple tasks and multiple targets.

If Executive Attention is merely a suite of strategies implemented in the cognitive processes that constitute the cognitive executive, there is no problem with it being so divided in a cognitive economy, since different particular executive processes can all implement it simultaneously in parallel without anything to unite them into a single indivisible whole. The (likely) fact that phenomenal experience is not divisible in this way suggests that it may not simply be a feature of a cognitive mechanism in the way that attention is, and that it might not even be of the same ontological class as attention. In fact, phenomenal unity is one of the things that makes consciousness so unique, and exploring the relationship between phenomenal FOC and Executively Attended higher-order content is a promising avenue for understanding why.²⁷¹

9.5.2 Theories of Consciousness

The holy grail of contemporary philosophy of mind is arguably a convincing theory of consciousness. The lack of such a theory is reflected in the broadness of the options that are still currently considered to be viable.²⁷² If there can indeed be $C \sim A$ and $A \cup C$ is the Scenario that obtains—if Executive Attention is not necessary or sufficient for phenomenal consciousness—then any successful theory of consciousness must account for this. While there is so much more that can be said, in this section I content myself with some brief observations on how my conclusions might be applied to Higher Order theories of consciousness, consequences for the Function Question as it relates to consciousness, and some methodological considerations.

²⁷¹ I explore this possible ontological uniqueness of consciousness—"Consciousness Uniquism"—in a paper currently under preparation.

²⁷² Chalmers' (2010) taxonomy of theories of consciousness is quite a useful one. See also van Gulick (2017).

9.5.2.1 *Higher Order Theories*

Some theories of consciousness are incompatible with my WHM of how content becomes conscious. Most obviously, Higher Order theories of consciousness (Carruthers, 2016) hold some sort of recursive processing to be necessary for phenomenal content, such that all phenomenal content arises via Executive Attention by definition.²⁷³ My arguments in Chapters 7 and 8 in favour of phenomenal overflow therefore have the effect of casting doubt on Higher Order theories. To be clear: they cast doubt on the *necessity* of higher-order cognition for phenomenality, not on its ability to produce phenomenality. My arguments do not discount phenomenality arising from higher-order cognition—they discount the idea that phenomenality can arise *only* from higher-order cognition.²⁷⁴

Forman (1999, p. 626) points out that genuine cases of what I have called global Pure Consciousness (6.3.3)—phenomenality in the absence of any content—would count strongly against theories of consciousness that merely identify consciousness with particular cognitive functions (functionalism) or see consciousness as a side effect of particular cognitive functions (epiphenomenalism), further dissociating consciousness from Executive Attention. My analysis suggests that while this kind of argument may be valid, its most important premise—genuine cases of global Pure Consciousness—is diabolically difficult to establish. However, foveal phenomenal overflow, which I have argued is on much more solid ground, may be enough to cast doubt that consciousness can only arise from Executive Attentional activity in a Higher Order theory. But of course, my analysis of no-report paradigms and neural signatures in 7.4.3 is grounds for pessimism that we can ever empirically resolve this question.

²⁷³ E.g., Rosenthal (2011, p. 431) “If somebody is in a mental state but doesn’t seem subjectively to be in that state, the state is not conscious.”

²⁷⁴ My account can embrace a “modest” higher-order theory, but not an “immodest” one (Block, 2011b, p. 421). For further arguments against ambitious higher-order theories, see Farrell (2018).

9.5.2.2 *The Function Question*

It has been suggested that a theory of consciousness cannot be divorced from an understanding of the function of consciousness, and that consciousness cannot be investigated apart from function (Cohen & Dennett, 2011). The function of attention is relatively uncontroversial. It was implicit in my definition of attention as a *suite of strategies for structuring cognition for the further processing of content*. Attention facilitates, perhaps even makes possible, the further processing of content. That is the functional role it plays. As such, Liberal Attention is involved in virtually every cognitive function, and Executive Attention is involved in every executive function. Attention therefore has deep roots throughout the cognitive economy.

The functional role of consciousness, on the other hand, is far more difficult to identify (Tsuchiya et al., 2015, p. 757). Debate continues over whether consciousness *per se* is causally efficacious or merely epiphenomenal (Kriegel, 2004). Some have argued that consciousness is important in the integration of information (Baars, 1988; Tononi, 2008), or for volition (Pierson & Trout, 2017). However, once you have subtracted out processes that are known to be capable of occurring unconsciously, there seems nothing left for consciousness *per se* to do in terms of integration or even volition. Similar subtractions can be made for other proposed functional roles such as rational thought, control of actions, etc., even the most complex of which can in the right circumstances (e.g., with much rehearsal) become subconscious. Indeed, Hassin (2013) has proposed that virtually any cognition is capable of proceeding without consciousness, an hypothesis that seems likely to be confirmed by the rapid advance of artificial intelligence (9.5.4.1).

This opens up the controversial possibility that consciousness itself may not have an integral role to play, but merely be an evolutionary “spandrel” (Haladjian & Montemayor, 2015; Robinson, Maley, & Piccinini, 2015)—a chance by-product of other cognitive developments that came along for a free ride. More specifically, though, phenomenal FOC that is never Executively Attended—like Ben’s overflowing speckles—seems especially profligate. What is the point of devoting cognitive

resources to producing so much rich content and making it phenomenal when it contributes nothing at all to our ongoing practical lives? We can only speculate at this stage, but I venture some thoughts.

First, the need for this profligate phenomenal content may be related to our human need for other apparently profligate pursuits such as beauty, art, music, and meaning, all of which may be enhanced by the richness of visual and other sensory experience. These things do not contribute *directly* to our survival or reproductive potential (the drivers of evolution) but they do enrich our lives and constitute a large part of what it means to be human.

Second, and much more prosaically, such needs may contribute *indirectly* to evolutionary fitness. For example, a capacity for powerful affective appreciation for beauty may be conducive to the choice of healthy reproductive mates (De Ridder & Vanneste, 2013; Grammer, Fink, Møller, & Thornhill, 2003). While phenomenally overflowing foveal vision itself seems unlikely to contribute to this, it may be useful in a third way: having all that rich phenomenal content *available* provides more options for Executive Attention, thus enhancing its effectiveness. Ben may not usually count all the speckles he phenomenally experiences, but if ever his life (or reproductive success) depended on it, his overflowing content provides him with the conscious knowledge that there are a lot of speckles there (Seeming-that) and the potential to count them one-by-one or in small subitised groups by shifting his Executive Attention rapidly and serially among them. As I say, all this is quite speculative, but worthy of further investigation. Whatever the case, if phenomenal foveal overflow occurs, it is an important datum in discussions about the function of consciousness, and may therefore inform discussions about theories of consciousness.

9.5.2.3 *Methodological Implications*

The temptation to use attention as a proxy for studying consciousness is one that has beguiled many (Allport, 1980, p. 113; Mole, 2008a, pp. 87–88; Posner, 1994), although it has also been challenged (Hardcastle, 1997; Lamme, 2003). If my analysis is correct, then this is a temptation we should avoid. Not only do attention and consciousness occur independently of each other, but the kind of relationship (4.4) between them remains unclear. I concluded above that IH and MC are highly unlikely, on the evidence available. This left one form of PC and many varieties of CA as the remaining possibilities. CA entails the possibility of a complex set of sub-options. Using attention to understand consciousness (or vice versa) depends on identifying and understanding the structure of the causal relationship between them, before we can draw any inferences from this approach with any kind of certainty. That seems premature at this stage. Nonetheless, these sub-options provide hypotheses that deserve investigation.

Although it is dangerous to use attention as a *proxy* for consciousness, there are valid uses of it as a *marker* of consciousness. Ben's confident Seeming—that is the marker of his rich phenomenal FOC (7.4.2), although its content is different. There is broad agreement that many cognitions instantiate the A&C Combination, and in those cases—depending on the kind of relationship between them—there may be a strong case for understanding consciousness through studying attention in certain ways. For example, Bayne (2013) makes a cogent case for *agency*—which plausibly depends on attentional strategies—as a marker of consciousness. My point is that this should never be taken as given, but always considered carefully. Certainly, if Ben's Executively un-Attended speckles are phenomenal, his phenomenal experience of them is a good place to explore the basis of consciousness, stripped bare of as much detritus as possible.

Finally, I take the pessimistic conclusion that no-report paradigms can never be successful at identifying $C \sim A$ in 7.4.3.3 as a challenge to seek new and creative ways to get around this serious epistemic obstacle. The central problem was that to calibrate something as a signature of consciousness, we inevitably rely on report in

the first instance, which means that we can never be sure that that signature is not just a signature of report (which involves Executive Attention) rather than a signature of consciousness *per se*. While my intuition is that this obstacle is ultimately impenetrable, any progress made will likely require highly creative thinking that accepts the intrinsically first-person nature of phenomenal experience and leverages it to good effect, even if that means we must rethink what constitutes the domain of the science of consciousness (*pace* Cohen & Dennett (2011), see 9.5.2.2).²⁷⁵

9.5.3 Other Cognitions

In 5.3.2.4 I posed the Integration Question, which is about how attention, consciousness and WM integrate together and with the rest of cognition. In this section I focus on two other cognitive concepts—intelligence and symbolic cognition—and sketch some reflections and applications in the light of my arguments and conclusions.

9.5.3.1 *Intelligence*

The triangulation of WM with attention and consciousness has an interesting application in the study of intelligence. There has been much work positively correlating individual differences in WM capacity and Executive Attention with *Gf* or

²⁷⁵ Dennett's (2003, 2007) *heterophenomenology* accepts first-person phenomenology but only as a datum in need of third-person verification. I take my arguments for immediacy to provide verification of a kind, but not, I suspect, the kind that will satisfy Dennett and Cohen. For a contrary view, see Kriegel (2007). Space prevents further exploration of this issue.

fluid intelligence (Kane & Engle, 2002, p. 638). Against this, it might be thought that some empirical measures of attention do not seem to correlate so well to intelligence. For example, subjects with attention deficit hyperactivity disorder (ADHD) show a broad spread of IQ scores, including very high intelligence (Katusic et al., 2011). However, as discussed briefly in 6.3.3.1, ADHD is not a diminution in attentional activity, but merely an inability to maintain attention to a single target over a period of time. The ADHD subject pays attention just as much as anyone else, just not to the same thing over time. ADHD is thus no exception to the correlation between Executive Attention and intelligence.

Does phenomenal consciousness play a role in intelligence? I mentioned the ongoing debate over whether phenomenal conscious plays a functional role in the kinds of complex cognitions that distinguish humans from other animals above (9.5.2.2). Kane and Engle (2002, p. 637) note that damage to the prefrontal cortex (PFC) results in deficits to core Executive Attention and WM functions, including, among other things, “attention, motor control, spatial orientation, short-term memory, temporal and source memory, metamemory, associative learning, creativity, perseveration, and reasoning.” Yet in none of these cases is it thought that there is any kind of deficit of phenomenal consciousness. The participation of PFC-damaged subjects in laboratory testing requires that they be physiologically conscious, which so far as we can tell is an indication of their phenomenal consciousness as well. This not only suggests that phenomenal consciousness plays no necessary role in intelligence, but also that PFC activity may not be necessary for phenomenal consciousness, as Lamme (2003) and others have argued, and as my WHM suggests.²⁷⁶ Indeed, we call computers “intelligent,” not because we think they are conscious, but because they are capable of complex functions, functions that rely upon the implementation of strategies of Liberal and Executive Attention (see 9.5.4.1). Attention, then, is intimately connected to intelligence in ways that consciousness is not. Empirically, measures of intelligence rely on measures of attentional abilities that have little to do with consciousness.

²⁷⁶ It is at best suggestive, since of course, significant portions of the PFC remain active.

9.5.3.2 *Symbolic Cognition*

A specific example of human intelligence is our unique ability to understand, manipulate, and creatively produce symbolic content: e.g., language and mathematics. Q plays an interesting role in a particular aspect of symbolic cognition—so-called 4E (embedded, embodied, enactive, and extended) cognition (Menary, 2010). In this section I lay out the evidence that attention and consciousness can dissociate for complex symbolic cognition in the brain, and consider some possible applications of the dissociation between attention and consciousness to the extended cognition discourse.

The connections between symbolic cognition and Executive Attention in WM are obvious. But it was thought that complex symbolic cognition requires conscious processing (Deutsch, Gawronski, & Strack, 2006; Dijksterhuis & Nordgren, 2006). Recently, however, evidence has been found for complex rule-based symbolic cognitive processing in the absence of phenomenal consciousness. Examples include masked simple addition instructions (Ric & Muller, 2012), semantic processing and multistep arithmetical calculations under continuous flash suppression masking (Sklar et al., 2012),²⁷⁷ implicit sequence learning, where the symbols are conscious but the learning of relations between them is subconscious (Mudrik, Faivre, & Koch, 2014), and sophisticated syntactic processing even in the absence of semantics under continuous flash suppression (Hung & Hsieh, 2015).²⁷⁸

All these cases constitute $A \sim C$, but things may not be so straightforward. There is an emerging pattern that suggests that consciousness is required for the integration of content in certain conditions:

²⁷⁷ But see Moors and Hesselmann (2018) and Rabagliati et al., (2018) which cast some doubt on this paper's conclusions, and on the broader landscape of complex non-conscious cognition.

²⁷⁸ See also a longer list of studies cited in Gelbard-Sagiv et al., (2016, pp. 1–2).

“Review of existing empirical data (especially in the visual domain, because most studies focus on visual awareness) suggests that there is no absolute dependency of integration on consciousness. Rather, the more complex or novel the stimuli, the more likely consciousness will be needed for integration to occur” (Mudrik et al., 2014, p. 488).

Complexity is not necessarily tied to consciousness. Expert, over-rehearsed actions of great complexity are performed without the subject’s conscious control or conscious perception of fine detail (Christensen, Sutton, & McIlwain, 2016). It seems that it is the *novelty* that most commonly requires the involvement of consciousness. The same is true of complex symbolic cognition.

Symbolic cognition quite often involves external media, which greatly expands human abilities. Take the example of solving a complex mathematical problem algebraically, one that needs two or three pages of multi-term equations. The capacity limitations of WM prevent the mathematician from holding anything like the complex series of equations and performing the manipulations in her head. At the least, there are huge gains in both efficiency and accuracy when she writes her working down. Literacy also involves not only semantic processing in the brain, but conducting conversations, recording words on a page or a computer, sharing them with others, etc., (Oatley & Djikic, 2008). This interplay between what is happening in the brain and what is happening beyond the brain has led some to conclude that “cognitive abilities are neither solely, or essentially, neural” (Menary, 2007, p. 622).

These are very deep waters, and I can only here point to some possible ways my approach to defining attention might help illuminate some of their dark depths. On my Operational definitions of attention as a suite of multiply-realizable strategies, there is no reason to think that only neuronal systems are capable of implementing them. Indeed, I will argue below (9.5.4.1) that machines implement them in quite complex ways. However, that is no reason to *thereby* ascribe consciousness to the

mechanisms that implement this attention, since attention is not an infallible proxy for consciousness (9.5.2.3).

There is an interesting question about where one might draw the line, then, between what constitutes human cognition, and what constitutes nothing more than an aid to human cognition. If one is happy to consider the processing carried out by a supercomputer as a kind of cognition, then what is to stop us considering the lines of math on the page as part of the mathematician's cognition? If one stipulates consciousness as a mark of cognition, the cases of unconscious complex symbolic cognition listed above make that untenable. Indeed, there are reasons to think that even within the brain, cognitive operations are not just conducted by the neurones, but require the active involvement of the glial structures in which neurones are embedded.²⁷⁹ If glial cells—whose activities directly contribute to the implementation of Operational Attentional strategies—can be said to participate in cognition, why not a calculator? There are many objections that can be raised here. I raise these questions only as pointers to interesting future applications of my approach to these kinds of questions.

9.5.4 Ethical Applications

The observation that computers require attention for intelligence but not consciousness points to some interesting ethical applications of the dissociation between attention and consciousness. In this section, I consider three such areas: artificial intelligence; animal ethics; and moral responsibility.

²⁷⁹ The neglected field of extra-neuronal involvement in cognition is a fascinating one (De-Miguel & Fuxe, 2012; Fields, 2009, 2013; Guidolin, Albertin, Guescini, Fuxe, & Agnati, 2011; Syková & Vargová, 2008).

9.5.4.1 *Artificial Intelligence*

The recent advances in artificial intelligence (AI) systems with prodigious capabilities pose two problems on which my discussion of Q might have some bearing. I argue that if we wish to address the dangers of rapidly advancing AI, it is attention that is most relevant, while if we wish to address questions of the value and rights of highly advanced AI, it is consciousness that matters.

The fear of machines becoming sentient and taking over the world of humans is both ancient and modern.²⁸⁰ One ancient form to be found in many cultures involved the need to keep the “barely human” slaves from rising up and harming their aristocratic masters. In modern times, this fear has come to be applied to non-biological slaves—machines—and may take a number of forms. One form is the fear over humans becoming superfluous. As a high school student in the 1970s, I vividly recall my Economics teacher gushing about how increasing automation will free humans for a life of leisure and pleasure. His predictions were dead wrong. Automation takes employment away from humans and leaves them not in pleasure, but in poverty (Ford, 2013). Another fear is that machines will be untrustworthy, liable to serious, perhaps even fatal mistakes (Parnas, 2017). And of course, science fiction is replete with tales of superintelligent machines rising up against their intellectually inferior human masters and enslaving them, or worse, eliminating them altogether.²⁸¹ More realistic scenarios involving the exponential growth in computing power and the emerging ability of computers to self-learn and self-correct have also been proposed, such as the singularity (Dubhashi & Lappin, 2017).

²⁸⁰ For some less dramatic questions, see Floridi et al., (2018).

²⁸¹ The Borg of *Star Trek: The Next Generation*, the Cylons of *Battlestar Galactica*, and the Replicators of *Stargate SG1* spring to mind as examples, not to mention the anxiety underlying Isaac Asimov’s Three Laws of Robotics, all of which are intended to protect humans from any potential malicious intent of their robotic slaves. A recent and more focused exploration of moral issues arising from human-like machines is the AMC/Channel 4/Kudos produced television series, *Humans* (2015-, <https://www.imdb.com/title/tt4122068/>).

On the other hand, as AI advances, it can become more and more human-like, at least in outward behaviour.²⁸² In one episode of the series *Star Trek: The Next Generation*,²⁸³ the question is posed in a particular way: a scientist wishes to disassemble an android named Data in order to better understand how he works, in the hopes of learning enough to reproduce him, for the benefit and service of humanity. The scientist bases his right to pursue this project on the grounds that an android like Data is merely a machine. No one would object to pulling apart the ship's on-board computer for a noble reason, nor would anyone even think to ask its permission to do so, since as a machine, it is mere property. Data is just an elaborate computer, so why would anyone treat "it" any differently?

But Data is not hidden behind a glass screen and keyboard. He is an android, which means he *looks* human. Like the ship's computer, he carries on conversations, but his conversations are far more intricate and humanlike. Further, he interacts physically with humans, shaking their hands, looking them in the eye with his own humanlike eyes, helping them out of tight squeezes by applying his formidable problem-solving skills and physical strength. Data's captain, Picard, after initially being ambivalent or even sympathetic to the disassembly project, given its potential benefits for humanity, eventually comes to defend Data's right to be treated not as property, but as a person. One of the arguments that sways him is the similarity between a future multitude of androids being treated as property and the past multitude of human slaves who were treated as property. Until we better understand just what Data is, Picard argues, we cannot run the risk of enslaving multitudes of possible persons.

Such questions are not purely academic, nor are they confined to fiction. recently, Google revealed a new feature in its Assistant that allows it to ring people on your

²⁸² There is little doubt that brains process information in different ways to computers. An interesting question I cannot go into here is whether non-human systems need necessarily to use the same attentional strategies I catalogued in my Taxonomy (3.3.4), or whether they might achieve the same ends by implementing different strategies, or perhaps the same strategies in different ways.

²⁸³ Series 2, episode 9, *The Measure of a Man*, first aired February 13, 1989.

behalf and conduct a conversation that would pass the Turing test (be indistinguishable from an all-human conversation).²⁸⁴ As machines approximate human behaviour more and more closely, our view of them is also likely to be drawn closer to our view of actual humans.

My analysis of attention and consciousness would suggest that the two concepts play different roles in discourses such as these. The *dangers* of AI do not depend on the AI becoming conscious, but on practical abilities. An AI program that determines that human beings are bad for the environment and must therefore be eliminated in order to preserve all other species need not be conscious to come to such a logical conclusion. In fact, it is all the more chillingly likely to come to this conclusion if it does not possess the kind of conscious ability for sympathy we normally associate with human experience. What will be critical in this kind of situation will be the effectiveness of the AI's processing, which in turn requires attentional strategies to be implemented.

On the other hand, when we come to determine the *value* to be placed on maintaining the integrity or even existence of an android like Data, or when we come to decide whether such an entity has *rights*, it is not attention we think about. It is not mere complexity of computation, behaviour, or ability that warrants our moral respect—otherwise tracking security cameras and supercomputers would be considered persons. Rather, I believe it is a very complex set of factors, among which computing ability is very minor indeed. Much more influential are very subjective (and often interrelated) factors such as the formation of an emotional bond with the entity, similarity in important ways to ourselves, and membership of our own “tribe” or “group” in some sense. Consciousness plays an important role in these kinds of factors, since we are ourselves essentially conscious beings. Data's case is ambiguous precisely because of his similarity—externally viewed at least—to human beings. The somewhat poorly defined question of “sentience” is raised as a

²⁸⁴ “Google I/O 2018: Assistant impersonates humans to make phone bookings” Adam Turner, *Sydney Morning Herald*, 9 May 2018, <https://www.smh.com.au/technology/google-i-o-2018-assistant-impersonates-humans-to-make-phone-bookings-20180509-p4ze52.html> accessed 15 May 2018.

major determinant in the proceedings in the episode—if it is even possible that Data is sentient, then he should be accorded the same value and rights as a human being.²⁸⁵ It is not cognitive abilities like attention that are relevant to ascriptions of moral value, but phenomenal consciousness.

A good rule of thumb, then, might be that if we are worried about machines taking over, it is attention and not consciousness that we should be concerned with, but if we are worried that we are treating our machines ethically, it is consciousness, not attention that we should be concerned with. I argued above (9.5.2.3) that in humans we cannot use attention as a proxy for consciousness. It would be interesting to explore whether a case for a Very Strong Reading of Q could be made, on which attention would not be reliable indicator of the presence of consciousness in machines as well as humans. One major issue is that machines are not “mushy” like human brains and almost certainly operate on different principles. Just what machine consciousness would be like, and what it would take to produce it, are huge questions that are largely opaque to us at present.

9.5.4.2 *Animal Ethics*

A similar issue arises in the field of animal ethics. For example, Singer (1975) and many who have followed him argue that animal interests ought to be respected because animals are capable of experiencing *suffering*—a concept that is inherently phenomenal.²⁸⁶ His concept of speciesism—discrimination on the grounds that an individual belongs to a different species—is reflected in my argument above that we ascribe moral value chiefly on the basis of similarity, in this case, belonging to the

²⁸⁵ Wisely, it is left up in the air as to whether or not Data is sentient. The Problem of Other Minds suggests we can never be certain whether even other human beings are sentient.

²⁸⁶ That animals are conscious appears to be a consensus position among academics (*Cambridge Declaration on Consciousness*, 2012). See also Andrews (2016), Godfrey-Smith (2016), and Barron and Klein (2016).

same species. But our empathy with even a non-human organism is mediated by the thing we have in common: the *conscious* experience of pain. This empathy may be what makes us ascribe more moral value to an animal than to a much more intelligent machine.

The other big moral question often discussed in *responsibility*. On what basis do we ascribe moral responsibility to an entity? We hold healthy adult humans morally responsible for their actions, but generally, not machines, animals, or small children. What light can attention and consciousness shed on this question? These too turn out to be deep waters.

Consider one account of moral responsibility: reasons responsiveness (Kennett & Fine, 2009, p. 86). On this account, an entity may be held morally responsible only insofar as it is capable of responding appropriately to reasons. Reasons responsiveness is not an inherently phenomenal process, but it is inherently attentional (Bello & Bridewell, 2017): appropriate responses are the outputs of processes that *select* the right inputs for *further processing* to produce a *coherent* output and *control* behaviour. The interesting thing is that machines, animals, and human infants are all quite capable of doing all of that. Yet we generally do not hold them morally responsible for their behaviour. In the case of machines, we transfer the responsibility to the human being who programmed or operated it. The machine just follows orders and has no choice in the matter. Similarly, we hold animals and human infants to be the slaves of their inbuilt instincts, and therefore not morally responsible. In both cases, it turns out that it is not the faculty for reasons responsiveness as such that underwrites the ascription of moral responsibility, but the faculty for control over how attention is implemented and executed, something more like the “ability to do otherwise” (Frankfurt, 2003).

The ability to do otherwise is a controversial concept, but one that is on some accounts tied to consciousness. A machine executing its programming “blindly” is not held morally responsible, whereas an adult human capable of reflecting on her behaviour and changing it in response, is. The animal and the infant fail to be morally responsible, not because they lack consciousness (they don't), but because

they lack the mature conscious *reflection* of the adult human. Moral responsibility, then, appears to be ascribed on the basis of properly functioning *phenomenal* higher-order (i.e., Executively Attended) processing—Experience-of, Seeming-that, and others. It is not ascribed if either the phenomenality is missing, or the Executive Attention is missing.

We can update our rule of thumb thus: Executive Attention is relevant to the dangers of AI; phenomenal consciousness of any kind is relevant to the ascription of moral value; but it is only phenomenal Executive Attention of a particular kind (reflective reasons responsiveness) that is relevant to the ascription of moral responsibility. This rule of thumb will, however, be challenged by machines that are capable of learning and developing themselves, independently of their human programmers. That is a fascinating topic I sadly do not have the space to pursue any further here.

9.5.4.3 *Moral Responsibility*

Considering machines (9.5.4.1) suggests that consciousness is what is important in ascriptions of moral *value* to entities, while considering animals and children suggests that conscious attention of the right kind is what is important in ascriptions of moral *responsibility* (9.5.4.2). In this section I consider whether cases of $A \sim C$ and $C \sim A$ in healthy adult humans are amenable to the ascription of moral responsibility.

Consider first an action that arises via $A \sim C$. We praise the sportsperson for their exquisite reflex shot, yet we hold a murderer blameless for a crime committed while sleep-walking (Ohayon & Schenck, 2010). Both are cases of $A \sim C$, yet we respond to them in opposite ways. Why? First, there is a question as to whether we *ought* to respond to them in different ways. A case can be made that in the same way that the sportsperson practiced voluntarily and consciously for years to be able to make that $A \sim C$ reflex shot, the sleep-walker too voluntarily and consciously allowed violence and hatred to fester within for years, resulting in the horrible violent act the

moment that normal societal controls were disabled in sleep-walking. But this is a weak argument. A person's responsibility must be taken as a whole, there must be a "unity of agency" (Kennett & Matthews, 2014). All of us have good and bad intentions within, but we are held responsible for the ultimate choices we make. The sportsperson's expertise and even the reflex shot itself are the ultimate result of the sum total of her voluntary, conscious decisions. She could have done otherwise (not practice, not care about winning, etc.) But the sleep-walker's crime is the result of one aspect of his cognitive life escaping the rest of his cognitive control in very unusual circumstances. Conscious, voluntary control can be lost in sleep-walking (Cartwright, 2004). We respond in opposite ways because these are not equivalent cases of $A \sim C$. They pose no threat to the principle that conscious attention of the right sort is necessary for the ascription of moral responsibility.

What of $C \sim A$? This treatise has suggested that if it is to be found anywhere, it will be in local phenomenal FOC, bereft of Executive Attention. I illustrated this in the case of phenomenal visual FOC, but of course, similar cases may be made for other modalities, perhaps (though this is less certain) even for subsets of the contents of memories, imaginings, and actions. In any case, the lack of Executive Attention would seem necessarily to indicate a lack of the kind of reflection—conscious or otherwise—that results in an ability to do otherwise (Bello & Bridewell, 2017). The healthy adult human, then, cannot be held morally responsible for cognitions that are $C \sim A$, or for their consequences, unless of course, conscious Executive Attention is subsequently applied to produce those consequences.

9.6 Final Reflections

I began this quest with a serene stroll through a lush rainforest, and ended with machines taking over the world, much like a blockbuster movie. But this journey has led us through far more interesting terrain—the mysterious inner cosmos of the human mind, that “glory, jest, and riddle of the world.”²⁸⁷ I said at the outset that this research was motivated by sheer unbridled and unapologetic curiosity about the mind, and especially about the enigma that is consciousness. If we are to improve our understanding of consciousness, one of our approaches will be to strike it against other aspects of cognition and see what sparks fly from the encounter. That has been the approach of this treatise. But there are many more ways to approach the question.

There is so much more that can be said about Q. I conducted this research with as open a mind as possible, and indeed, changed my mind about the answer to Q a number of times along the way. Perhaps this treatise is best viewed as a snapshot capturing a single frame in a movie with a long past and (hopefully) a long future to come. There is no telling where it will end, but the joy is in the journey. It is a privilege to be part—however small—of that wonder-filled and very human quest to understand what it is that makes us who we are.

²⁸⁷ Alexander Pope, *An Essay on Man*, 2.1.

Appendices

Appendix 1. Distinctions between different kinds of attention.

In this appendix I flesh out the distinctions briefly summarised in 3.2.3.

Internal v External

The location of the object of one's attention provides the first of our distinctions, that between *internal* and *external* attention (Buckner, Andrews-Hanna, & Schacter, 2008).²⁸⁸ Internal attention is directed to the subject's own internal mental states (e.g., a sense of general sadness, or the number six), whereas external attention is directed to objects in the environment (e.g., a mosquito buzzing near one's ear).

Endogenous v Exogenous

A related distinction is that between attention that is *top-down* or *endogenous*, or *bottom-up* or *exogenous*.²⁸⁹ This is a distinction that deals with how attention is recruited. Top-down attention is directed by the subject's own internal, higher goals, as when you are purposefully looking for a word in a puzzle. Or attention may be captured bottom-up by external, salient stimuli, as when you automatically turn

²⁸⁸ Fox et al's (2005) distinction between *task-negative* and *task-positive* brain networks may be related to this distinction, although there are some questions that might be raised.

²⁸⁹ Other authors employ different terms for this distinction, e.g., *goal-directed* (endogenous) v *stimulus-driven* (exogenous), (Corbetta & Shulman, 2002).

your head in the direction of a loud bang. Desimone and Duncan (1995, p. 201) point out that sometimes, endogenous information—e.g., long term memory of the importance of certain perceptual features—can bias attention in bottom-up fashion, suggesting that these categories can interact in interesting and complex ways.

Voluntary v Involuntary

A closely related, though discrete distinction is that between *voluntary* or *involuntary* attention. This is a distinction that deals with whether the subject controls the direction of attention, or whether it is directed by factors outside the subject's voluntary control. Attention (again) may be directed by the subject's own higher goals, as when you are purposefully looking for a word in a puzzle. Or it may be captured without the subject's voluntary choice, and perhaps even against the subject's will, as when you can't get a song out of your head.

Some authors seem simply to identify top-down endogenous attention with voluntary attention and bottom-up exogenous attention with involuntary attention (Ciaramelli, Grady, & Moscovitch, 2008; Theeuwes, 1991). I believe this is a mistake. On the purely conceptual level, it is certainly possible to imagine, say, involuntary top-down attention—attention, perhaps, that is directed by an involuntary influence of priming. We should also avoid the temptation to think that this distinction simply maps neatly onto the conscious/non-conscious distinction. Voluntary does not necessarily mean conscious: the minutely detailed fine-motor adjustments that go into highly rehearsed actions like fine-tuning the angle at which one strikes a golf ball are in some sense voluntary, yet not conscious. Similarly, looking at a horror scene in a movie although one dearly wants to avoid nightmares arising from the sight—may be cases of attention that is in some sense involuntary, yet conscious. These examples are subject to challenge of course, and they touch upon difficult questions in the philosophy of the will that I need not delve into any further here, but I merely raise them to show that the assumption that voluntary equals conscious and involuntary equals unconscious need not be taken for granted.

Focal v Background

A fourth distinction—*focal* versus *background* attention—relates to the size of the field of attention. For example, one's available visual field is quite wide, but the section of the field to which one can pay attention may be relatively small. Braun and Sagi (1990, p. 45) cite evidence for the focus of visual attention subtending an angle of about 1° of eccentricity and Iwasaki (1993, pp. 213–214, 220) considers processing that occurs outside this “spotlight” of focal attention to be non-attentional processing. However, it is generally thought that the focus of attention may be variable, using the camera's zoom lens as an analogy for the ability to widen or tighten one's field of attention: “attention can vary from a uniform distribution over the entire field to a highly focused concentration” (Theeuwes, 1991, p. 83), although widening the attentional field seems to distribute the same processing resources more thinly over a larger area (C. W. Eriksen & St James, 1986). The border between focal and background attention may be a graded one, resembling the rather fuzzy umbra and penumbra of a solar eclipse (8.6.3). The border between focal or background attention and non-attention may also be graded or an on-off affair. We should resist the temptation to identify, *a priori*, focal attention with consciousness, and background attention with unconscious processes. The two distinctions are certainly conceptually discrete, and how they relate is a matter for empirical investigation.

Spatial, Feature, or Object

Fifth, the target of visual attention may be of three kinds: “defined by a circumscribed region in space (focal attention), by a particular feature (feature-based attention), or by an object (object-based attention)” (Koch & Tsuchiya, 2007, p. 16). I shall call these three *spatial*, *feature*, and *object* attention respectively. There are other ways of drawing this distinction. Duncan (1984) calls feature attention “discrimination-based attention,” while Desimone and Duncan's (1995) findings suggest to me that object attention may in a sense just be a special case of feature

attention, one where many features are bound together, rather than compared. Kanai, Tsuchiya, and Verstraten (2006, p. 2335) suggests that spatial and feature-based attentions at least have different neural correlates. They conclude that feature-based attention may be correlated sometimes with unconscious processing, while spatial attention may correlate with conscious processing, but again I caution against such identifications. While this distinction arose from work on visual attention, it may well have analogues in attention outside vision.

Synchronic v Diachronic

Sixth, *synchronic* views of attention describe it as a state of a cognitive process, a “snapshot” in time, whereas *diachronic* views of attention respect the fact that any process necessarily plays out over a period of time. There would seem to be no room for middle ground between these two poles. However one might ask whether we can speak of synchronic attention at all. It is common, for example, to speak of an “attention span,” the period of time over which one can maintain attention.

Attention may be deconstructed into smaller temporal components—Braun and Sagi (1990, p. 46) cite evidence for the duration of a single fixation of visual attention of between 20 and 100 milliseconds. And at least as far back as William James, the experience of the “specious present” has been understood as not being experienced as instantaneous anyhow, but rather as having some brief yet finite duration (Gallagher, 1998, Chapter 2; Le Poidevin, 2015), perhaps somewhere between half a second and three seconds (Dainton, 2000, pp. 170–171). Even the “now,” therefore, seems necessarily diachronic. And of course, attention can be over an extended period of time, as when one cares for the environment and attends to environmental issues over many decades.

We need not completely abandon the concept of synchronicity, though, so long as we understand these qualifications. The idea of a static “snapshot” is an incomplete picture of what attention really is, just as a photographic snapshot of an Olympic sprinter in full flight is an incomplete picture of what a race, or running, really is.

Personal/Subpersonal

Finally, attention may operate on a personal level (e.g., tracking a moving target with one's eyes and head), or on a subpersonal level (e.g. in early visual processing). The subpersonal, of course, admits of many levels. Personal and subpersonal attention are not mutually exclusive, and Watzl (2011b) considers the relationship between them. Some have argued that personal attention may just be reducible to subpersonal level attentional mechanism, while others have resisted this reductionism.

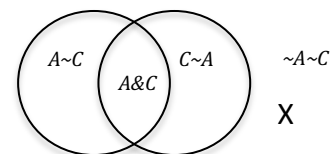
Appendix 2. Four Unlikely Scenarios

The four main Scenarios described in 4.2.3 each have an X in the $A \& C$ intersection, signifying that there are at least some cognitive processes that are both attended and phenomenally conscious—conscious attention. However, there is a group of Scenarios parallel to the four above, in which there is no X in the $A \& C$ intersection. These Scenarios do not figure in most discussions on Q. For these Scenarios, attention is neither necessary nor sufficient for consciousness, and consciousness is neither necessary nor sufficient for attention. This group of Scenarios is described below for completeness, even though it seems quite obvious to almost everyone that we often experience conscious attention.²⁹⁰

Table 5. Four unlikely STF Scenarios.

\emptyset (*Inattentive Unconsciousness*)

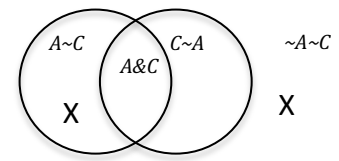
- There never is such a thing as either attention or consciousness. Thus trivially, there is no such thing as a process that is both attended and conscious.
- In set theory, " \emptyset " is the symbol for an *empty set* or *null set*, a set with zero members, which captures nicely the idea that no cognitive processes are ever attended or conscious.



²⁹⁰ Montemayor & Haladjian (2015, p. 4) cite Tallon-Baudry (2012) as favouring the view that there is no such thing as conscious attention.

$A \setminus C$ (Attentive Unconsciousness)

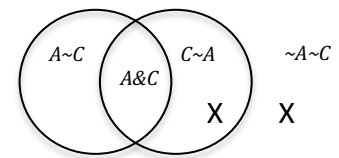
- There is only attention, but never is there such a thing as consciousness. Thus trivially, there is no such thing as a process that is both attended and conscious.



- In set theory, " $P \setminus Q$ " (*difference* of two sets:²⁹¹ members that belong to P but not to Q) captures nicely the idea that cognitive processes may be attended but cannot be conscious.

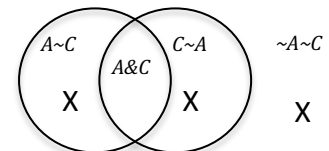
$C \setminus A$ (Conscious Inattentiveness)

- There is only consciousness, but never is there such a thing as attention. Thus trivially, there is no such thing as a process that is both attended and conscious.



$A \Delta C$ (Complete Independence)

- Some processes are conscious but not attended, some are attended but not conscious, but no processes are ever both attended and conscious, and of course, some are neither conscious nor attended. Attention and consciousness never occur together in the same process.
- In set theory, " $P \Delta Q$ " (*symmetric difference*: members that belong to P or to Q , but not to their intersection) captures nicely the idea that cognitive processes may be either attended or conscious, but not both.



²⁹¹ Also known as the *relative complement* of Q in P .

These four Scenarios are generally considered to be unlikely as a description of a normal human cognitive economy, although they are not completely impossible. They might obtain in certain unusual or imaginary circumstances.

The \emptyset Scenario describes a cognitive economy in which there are never any cognitive processes that are either attended or conscious. This sounds like a description of a deeply *comatose* patient in whom both attention and consciousness are impossible, although many unattended and unconscious cognitive processes are still proceeding, as evidenced by continuing activity on, say, the EEG trace.

Some extreme forms of *eliminativism* about consciousness (Churchland, 1983; Wilkes, 1984) might be taken as a case of the two Scenarios in which there is no consciousness ($A \setminus C$ and \emptyset). Less extreme forms, such as Dennett's (1991) scepticism about qualia, might also plausibly fall under these Scenarios. One might also read some forms of philosophical *Behaviourism* as being epistemic versions of the Scenarios in which there is no consciousness.

Perhaps an appropriately programmed *computer* falls under the $A \setminus C$ Scenario, since it may be capable of attending to stimuli (given the right definition of attention), and processing them, while most people would accept that a computer has no conscious experience of that processing.

An example of the $C \setminus A$ Scenario might be a pre-conception *Platonic soul*—a disembodied consciousness that is unattached to any brain, and therefore a pure consciousness without any kind of cognitive abilities, including attention. Another example might be some forms of *brain damage* that leave a subject with an inability to focus her attention, without impacting on her ability to have a conscious experience (although she may still be capable of a more diffuse kind of attention).

Finally, the $A \Delta C$ Scenario, in which both attention and consciousness are present in a cognitive economy without ever coinciding as properties of any single cognitive process, is a difficult Scenario to find examples of. Perhaps it might occur in some kind of hypothetical extreme form of brain damage that disconnects the two from

each other, such that, for example, one loses all visual consciousness but not visual attention, and one loses all auditory attention but not auditory consciousness. Another possibility is a computer with the ability to devote computing resources to certain input data over others, which is then “haunted” by a ghostly disembodied Platonic soul without access to its computations. I am aware of no actual situation in which this Scenario occurs.²⁹²

The concept of Liberal Attention (3.4.4) being ubiquitous in cognition would be captured by Scenarios that lack the Combinations $C \sim A$ and $\sim A \sim C$, namely the four asterisked Combinations: $A = C^*$; $A \supset C^*$; \emptyset^* ; and $A \setminus C^*$.

²⁹² Iwasaki (1993) seems to consider something like $A \Delta C$ seriously. Montemayor and Haladjian (2015, pp. 5–6) also include this possibility in their CAD framework.

Appendix 3. Tabulated Comparisons of STF with Other Frameworks

Iwasaki

Table 6. Comparison of STF with the descriptions of Iwasaki (1993, p. 212).

Iwasaki	STF
1. “people may perceive only those stimuli that are under focal attention”	$A = C$ (or $A \supset C?$)
2. “although attention may bring about some differences in what we perceive consciously, consciousness can exist independently of attention”	$C \supset A$ (or $A \cup C?$)
3. “consciousness is a mental faculty which works independently of attention. The distinctive feature of this view is its denial of any possibility of attentional modulation of consciousness”	$C \supset A$ (or $A \cup C?$)

Lamme

Table 7. Comparison of STF with the models of Lamme (2003).

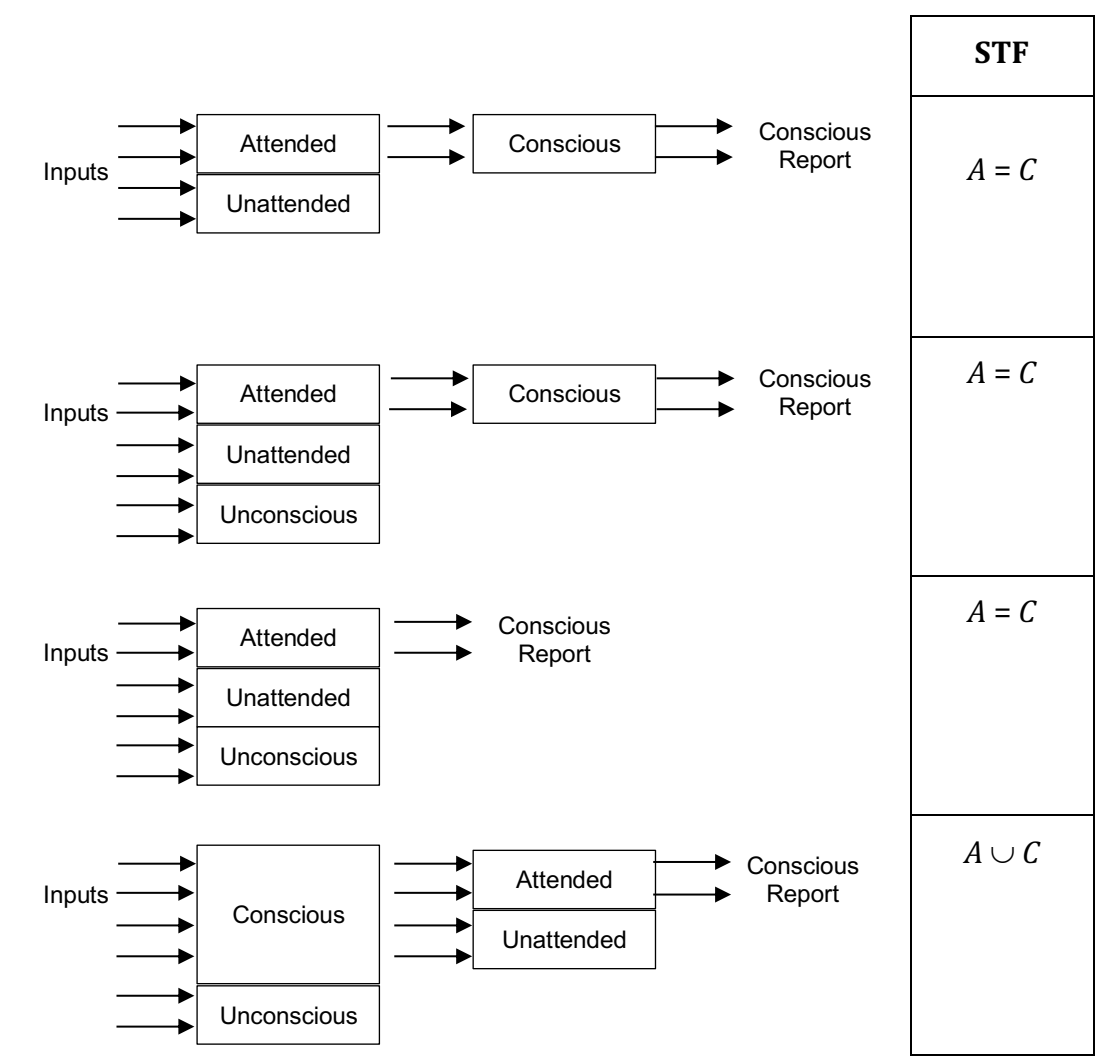


Figure 6 Four models of visual awareness and its relation to attention, modified from Lamme (2003) p. 13.

Schwitzgebel

Table 8. Comparison of STF with the categories of Schwitzgebel (2007).

Schwitzgebel	STF
<i>Rich Consciousness:</i> the contents of sensory experience at any moment in time are multimodal and full of details.	$C \supset A$
<i>Thin Consciousness:</i> the contents of sensory experience at any moment in time are limited to just a few modalities and details.	$A = C$ (possibly $A \supset C$)

Koch & Tsuchiya

Table 9. Comparison of STF with the framework of Koch and Tsuchiya (2007).

Fourfold taxonomy of conscious and unconscious percepts and behaviours of Koch and Tsuchiya (2007) p. 17, indicating the corresponding STF Scenario.

	Might not give rise to consciousness		Gives rise to consciousness	
Top-down attention is not required	Formation of afterimages Rapid vision (<120ms) Zombie behaviours	$\sim A \sim C$	Pop-out in search Iconic memory Gist Animal and gender detection in dual tasks Partial reportability	$C \sim A$
Top-down attention is required	Priming Adaptation Visual search Thoughts	$A \sim C$	Working memory Detection and discrimination of unexpected and unfamiliar stimuli Full reportability	$A \& C$

Dijksterhuis et al., (2010, p. 472) compare this taxonomy to the earlier more complex, but more general system of Wegner and Smart (1997).

De Brigard

Table 10. Comparison of STF with the propositions of De Brigard (2010, pp. 189–190).

De Brigard	STF
(a) Attention is necessary and sufficient for consciousness.	$A = C$
(b) Attention is necessary but not sufficient for consciousness.	$A \supset C$
(c) Attention is neither necessary nor sufficient for consciousness.	$A \cup C$
Mole (2008a) “believes that, according to our commonsense, while consciousness is necessary for attention, attention isn’t necessary for consciousness”	$C \supset A$

Van Boxtel

Comparison of STF with the framework of van Boxtel et al., (2010, p. 6, Figure 2a)

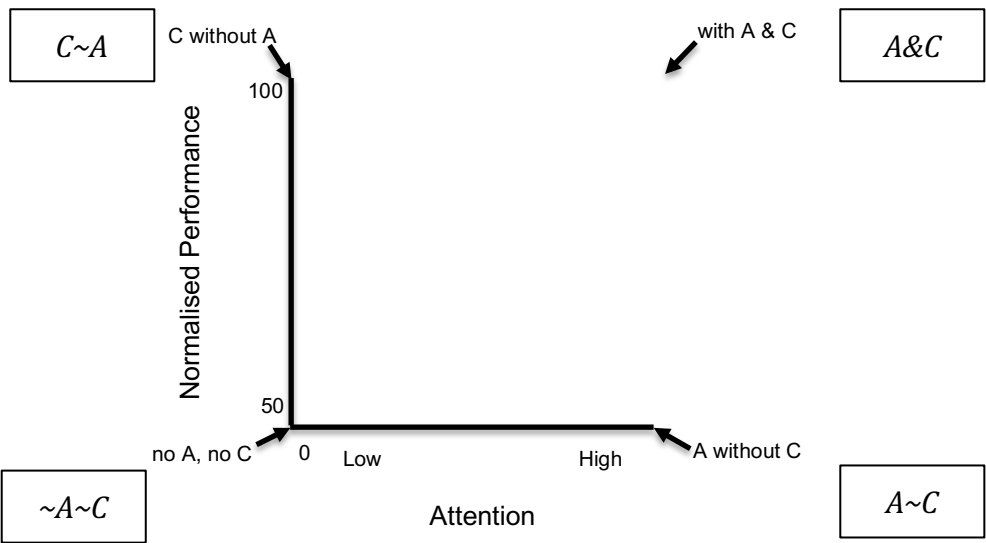


Figure 7 Consciousness vs Attention, modified from Van Boxtel et al., (2010, p. 6, Figure 2a).

Montemayor and Haladjian

Table 11. Comparison of STF with the CAD framework of Montemayor & Haladjian (2015, pp. 5–6).

CAD	STF
1. Identity. Consciousness is identical with attention.	$A = C$
2. Type-A CAD: <u>one</u> kind of conscious attention; attention is necessary for consciousness.	$A \supset C$
3. Type-B CAD: <u>one</u> kind of conscious attention; attention is not necessary for consciousness.	$A \cup C$
4. Type-C CAD: <u>many</u> forms of conscious attention, (thus increasing the opportunities for dissociation); attention is not necessary for consciousness.	$A \cup C$
5. Full Dissociation: A & C are completely separate in both description and mechanism; “there is no conscious attention.”	<i>Scenarios 5-8</i>

Pitts et al.

Table 12. Comparison of STF with the descriptions of Pitts et al., (2018, p. 2).

Pitts et al.	STF
“attention can modulate sensory processing even in the absence of conscious perception, and conscious experience can and does occur in the absence of attention”	$A \cup C$
“attention and consciousness are distinct psychological processes, and whereas attention can operate independently of consciousness, the reverse is not the case “	$A \supset C$
“A third view has also been proposed, which argues against any type of dissociation—i.e. consciousness <i>is</i> attention”	$A = C$

Appendix 4. Classification of Authors by STF Scenario

In this appendix I set out some methodological notes on how I derived the figures for Figure 3 (4.3.6) followed by a list of authors by Scenario (Table 13) and another of Scenarios by author (Table 14).

Some Methodological Notes.

- I make no claim whatsoever to this survey being scientific, and all observations drawn from it are to be taken as liable to revision.
- The sample set here is those texts I have consulted in the course of researching this thesis that directly address Q in sufficient detail to identify a particular Scenario.
- It is by no means exhaustive, and no doubt there are significant biases due to the research resources available to me, etc. Smaller sample sets might be derived by surveying seminal publications, such as the commentaries to Block's (2007) BBS article, the chapters of Mole, Smithies, and Wu (2011a), or the articles in the special issue of *Philosophical Transactions of the Royal Society B: Biological Sciences* (2018, volume 373, issue 1755) on "Perceptual consciousness and cognitive access."
- Many texts that address some narrow aspect of Q (e.g., whether or not a single Combination obtains) have not been included, since they do not provide sufficient evidence for a Scenario. In some cases, further research into the authors' other publications might identify a preferred Scenario, but as this survey is not the chief goal of my thesis, I have not pursued this project further.
- The ideas of a number of authors resist easy classification, and I have indicated this by "tentatively" classifying them. A few are so ambiguous that I left them without a Scenario ("Unable to Classify").
- Where a paper or set of papers have multiple authors, I have counted them as single authors (and used "et al."), unless they have independent papers that

differ on significant points (e.g., De Brigard and Prinz). This avoids a source of bias since philosophers tend to write alone, while scientists tend to write in groups, which would have heavily biased the counts in favour of the scientists.

- The academic fields of authors were determined by examining their departmental affiliations and/or resumes if publicly available.
- Where an author has qualifications in more than one field, the field of his highest-level qualification was preferred. Mostly, this corresponded with his departmental affiliation.
- I counted William James as a philosopher although as a founder of modern psychology, he also belongs under the non-philosopher (psychology) category.

Authors by Scenario

Table 13. Classification of Authors by Scenario.

[P] = philosopher—background primarily in philosophy.

[NP] = non-philosopher—background primarily in a field other than philosophy. See Table 14 for more detail on specific non-philosophical fields.

Authors	Citations	Comments
1. Scenario A = C		
William James [P]	(1890)	Read by Cowan (1999, p. 63) and many others as seeing attention as operating within consciousness. ²⁹³

²⁹³ But, see Hardcastle (1997).

Allport [NP]	(1980)	Read by Mole (2008a, p. 87) as using attention as “a code name for consciousness.”
Dennett [P]	(1991)	Read by Schwitzgebel (2007, p. 7) as espousing a thin view of consciousness, although he is co-author with Cohen of papers that take a $A \supset C$ view.
Stazicker [P]	(2011a)	Argues for a constitutive connection between attention and consciousness.
Baddeley [NP]	(Baddeley, 1993a, 2000)	Identifies consciousness with the activity of the central executive component of working memory (see Chapter 5).
Cowan [NP]	(1988, 1995)	Identifies consciousness with the activity of the focus of attention in working memory (see Chapter 5).
Hine [P]	(2010, 2015; 2010)	Consciousness can be explained by attention.
De Brigard and Prinz [P]	(De Brigard, 2010; De Brigard & Prinz, 2010; Prinz, 2011)	Attention is the gateway to working memory, consciousness arises from working memory.
Carruthers [P]	(2017)	Argues that Block’s phenomenal consciousness and access consciousness are not truly distinct. Conscious content is the

		content of the global workspace, though not of working memory.
O'Regan [NP] and Noë [P]	(2001)	Read by Cohen et al., (2012, p. 411) as $A = C$.
Posner [NP]	(1994)	Draws an analogy between the relationship between attention and consciousness and the relationship between DNA and "life."
Mack and Rock [NP]	(1998)	Pioneered the study of inattention blindness.

2. Scenario $A \supset C$		
Dehaene et al. [NP]	(Dehaene et al., 2006; Dehaene, Charles, King, & Marti, 2014; Dehaene, Lau, & Kouider, 2017, p. 489; Naccache, Blandin, & Dehaene, 2002)	Theorize as to the neural underpinnings of attention and consciousness, and propose a model that bears some similarities to that of Lamme (2003).
Cohen et al. [NP]	(2012)	Reject the suggestion that "attention and awareness can be doubly dissociated," yet affirm that "attention is necessary,

		though not sufficient, for conscious awareness” (p. 411).
Pitts et al. [NP]	(2018)	Argue for a “single dissociation view” (pp. 2-3) on which there can be $A \sim C$ but not $C \sim A$, albeit with some nuance and an open mind.
John G Taylor [NP]	(2002)	Asserts $A \sim C$ and denies $C \sim A$ (p. 206).
Richards [P]	(2013, 2016)	Critiques of Ned Block’s thesis that consciousness overflows attention.
Baars [NP]	(1988, 1997b)	Difficult to classify—but Global Workspace model is suggestive.
Soto & Silvanto [NP]	(Soto, Mäntylä, & Silvanto, 2011; Soto & Silvanto, 2014, 2016)	Tentative STF classification. Their “emerging framework” (2014) endows attention with the dual roles of bringing content into working memory and bringing working memory content into consciousness (a denial of $C \sim A$, compare De Brigard & Prinz) but they also argue strongly <i>for</i> non-conscious working memory ($A \sim C$).

3. Scenario $C \supset A$		
Smithies [P]	(2011)	“attention is essentially a phenomenon of consciousness” (p. 247)
Watzl [P]	(2010, 2011a)	Sees attention as a structuring of conscious experience, and argues that putative empirical cases of $A \sim C$ fail (2010, p. 322).
Wolfe [NP]	(1999b)	Introduces the concepts of <i>pre-attentive</i> and <i>post-attentive</i> consciousness.
Iwasaki [NP]	(1993)	Difficult to classify (4.2.4.2 and Appendix 3, Table 6).
Jennings [P]	(2012, 2015)	Affirms a particular type of $C \sim A$ she calls “conscious entrainment.”

4. Scenario $A \cup C$		
Block [P]	(2013b, p. 182; 2014, p. 556)	See Chapters 7 and 8 for detailed discussion.
Bronfman et al. [NP]	(2014; 2018)	Argue strongly for $C \sim A$. Assume $A \sim C$ in their 2018 (pp. 16-17) citing Kouider and Dehaene (2007) in support.

Armstrong [P]	(1995)	Accepts that Block's phenomenal and access consciousnesses can occur in isolation.
Hardcastle [P]	(1997)	Reinterprets William James to allow for $C \sim A$ (p. 56), and argues against the sufficiency of attention for consciousness (i.e., argues for $A \sim C$, p. 62)
Kentridge et al. [NP]	(Kentridge, 2011, pp. 233–235; Kentridge et al., 1999, 2004, 2008; L. J. Norman et al., 2015)	Empirical evidence for the dissociation of attention and consciousness in a broad variety of circumstances.
Crick et al. [NP]	(Crick & Koch, 1990; Koch & Tsuchiya, 2007; Tsuchiya & van Boxtel, 2010; van Boxtel et al., 2010).	KT controversially use the term "near-absence of attention" (4.2.4.2).
Lamme's [NP]	(2003, 2010)	Based on the neural activity patterns of <i>recurrent processing</i> and <i>feed forward sweep</i> .
Mole [P]	(2008a, 2011a, 2014)	<i>Cognitive unison</i> model of attention. In his 2014 (pp 43-49), he reverses his previous opposition to $A \sim C$.
Hohwy [P]	(2012)	Based on a predictive coding model of attention and consciousness, he affirms the

		possibility of all four Combinations on p. 5.
Montemayor [P] and Haladjian [NP]	(2015)	Discuss multiple processes of attention, consciousness, and conscious attention (see Appendix 3), dissociations that reflect $A\sim C$, $C\sim A$, and $A\&C$ respectively.
Dijksterhuis et al. [NP]	(2010)	Social scientists who accept the fourfold Combination structure of Koch and Tsuchiya (2007).
Hassin et al. [NP]	(2013; 2009)	Concept of implicit working memory is indubitably $A\sim C$. Affirms $C\sim A$ if one takes attention to be something like a spotlight amidst conscious content, but not if attention is just defined as the devotion of cognitive resources ²⁹⁴ (see 3.4.4);
Koivisto, Revonsuo, et al. [NP]	(Koivisto & Revonsuo, 2007; Koivisto, Revonsuo, & Salminen, 2005)	Electrophysiological evidence for the independence of awareness and attention.
Bayne [P]	(Bayne, 2010; Shea & Bayne, 2010)	Tentative classification based on suggestive statements in Shea & Bayne pp. 468 ($A\sim C$) and 469 ($C\sim A$).

²⁹⁴ Hassin, personal communication.

Tye [P]	(2010, 2014)	Despite arguing against Block's (2013b) argument for different grains of seeing and attention, he has himself argued for both $A \sim C$ and $C \sim A$
Tallon-Baudry [NP]	(2012, p. 7)	Argues for distinct neural bases of attention and consciousness, and explicitly says that her "cumulative influence" model suggests both $A \sim C$ and $C \sim A$.
Wegner & Smart [NP]	(1997)	Frame the Combinations in terms of <i>deep cognitive activation</i> and consciousness.
Wu [P]	(2014, pp. 150–172, 2016, p. 3)	Explicit and repeated affirmation of $A \sim C$. He tentatively affirms $C \sim A$ (2013, p. 1180), so a tentative STF classification of $A \cup C$.

Unable to Classify		
Barrett [P]	(2014)	Argues against Prinz's interpretation of the empirical evidence, but leaves the question open himself.
Phillips [P]	(2011a, 2011b, 2015, 2018)	Has argued against the $C \sim A$ of phenomenal overflow (Chapter 7) but maintains a sceptical agnosticism on Q.

John Henry Taylor [P]	(2013b, 2013a)	Argues against De Brigard and Prinz ($A = C$), but also against Block ($A \cup C$). His PhD thesis (2015b), which might clarify his position, is unfortunately under copyright embargo until 2020. I did not attempt to contact him for a copy.
-----------------------	----------------	---

Scenarios by Author

Table 14. Alphabetical list of Scenarios by author.

NOTES:

- “Other Disciplines”—I have included disciplines where authors have formal qualifications, as gleaned from their curriculum vitae. Sometimes, they are no longer actively engaged in these disciplines.
- Where an author’s discipline is more difficult to pin down, I have used the umbrella term “cognitive science” (Bechtel, Abrahamsen, & Graham, 2001; Simon, 1980).

Scenario	Name	Discipline	Other Disciplines
$A = C$	Allport, <i>Alan</i>	Psychology	
$A \cup C$	Armstrong, <i>David M</i>	Philosophy	
$A \supset C$	Baars, <i>Bernard</i>	Psychology	
$A = C$	Baddeley, <i>Alan</i>	Psychology	

Scenario	Name	Discipline	Other Disciplines
$?$	Barrett, <i>David</i>	Philosophy	
$A \cup C$	Bayne, <i>Tim</i>	Philosophy	
$A \cup C$	Block, <i>Ned</i>	Philosophy	
$A \cup C$	Bronfman, <i>Zohar</i> (et al.)	Computational Cognitive Neuroscience	<i>History and Philosophy of Science</i>
$A = C$	Carruthers, <i>Peter</i>	Philosophy	
$A \supset C$	Cohen, <i>Michael A</i>	<i>Cognitive Psychology</i>	<i>Philosophy</i>
$A = C$	Cowan, <i>Nelson</i>	Psychology	
$A \cup C$	Crick, <i>Francis</i>	Cognitive Science	<i>Physics Biology</i>
$A = C$	De Brigard, <i>Felipe</i>	Philosophy	
$A \supset C$	Dehaene, <i>Stanislas</i> (et al.)	Cognitive Science	<i>Cognitive Neuroscience, Applied Mathematics and Computer Science, Mathematics.</i>
$A = C$	Dennett, <i>Daniel</i>	Philosophy	
$A \cup C$	Dijksterhuis, <i>Ap</i> (et al.)	Social Science	<i>Social Psychology</i>
$A \cup C$	Haladjian, <i>Harry H</i>	Cognitive Psychology	
$A \cup C$	Hardcastle, <i>Valerie</i>	Philosophy	<i>Political Science Cognitive Science</i>
$A \cup C$	Hassin, <i>Ran</i> (et al.)	Cognitive Science	
$A = C$	Hine, <i>Rik</i>	Philosophy	
$A \cup C$	Hohwy, <i>Jakob</i>	Philosophy	
$C \supset A$	Iwasaki, <i>Syoichi</i>	Cognitive Science	

Scenario	Name	Discipline	Other Disciplines
$A = C$	James, William	Philosophy	<i>Medicine Psychology</i>
$C \supset A$	Jennings, Carolyn Dicey	Philosophy	<i>Psychology</i>
$A \cup C$	Kentridge, Robert W (et al.)	Psychology	
$A \cup C$	Koivisto, Mika (et al.)	Neuroscience	<i>Philosophy</i>
$A \cup C$	Lamme, Victor AF	Cognitive Science	<i>Medicine</i>
$A = C$	Mack, Arien & Rock, Irvin	Cognitive Science	
$A \cup C$	Mole, Christopher	Philosophy	<i>Psychology</i>
$A \cup C$	Montemayor, Carlos	Philosophy	
$A = C$	Noë, Alva	Philosophy	
$A = C$	O'Regan, John Kevin	Psychology	<i>Mathematical Physics</i>
?	Phillips, Ian	Philosophy	<i>Physics</i>
$A \supset C$	Pitts, Michael A (et al.)	Psychology	
$A = C$	Posner, Michael I	Psychology	<i>Physics</i>
$A = C$	Prinz, Jesse	Philosophy	
$A \supset C$	Richards, Bradley	Philosophy	<i>History and Philosophy of Science</i>
$A \supset C$	Soto, David & Silvanto, Juha	Cognitive Science	<i>Experimental Psychology</i>
$C \supset A$	Smithies, Declan	Philosophy	<i>Theology</i>
$A = C$	Stazicker, James David	Philosophy	<i>Classics</i>

Scenario	Name	Discipline	Other Disciplines
$A \cup C$	Tallon-Baudry, Catherine	Neuroscience	Biology
$A \supset C$	Taylor, John G	Mathematics	Physics
?	Taylor, John Henry	Philosophy	
$A \cup C$	Tye, Michael	Philosophy	Physics
$C \supset A$	Watzl, Sebastian	Philosophy	Biology
$A \cup C$	Wegner, Daniel M	Psychology	Physics
$C \supset A$	Wolfe, Jeremy M	Psychology	Ophthalmology Radiology
$A \cup C$	Wu, Wayne	Philosophy	Biology Chemistry

Appendix 5. Models of WM: A Selection

The name of a model is provided, generally by the authors, together with the chief originators of the model, and one or two central citations. This is followed in the second column by my summary of the model, and in the third column, some of my own thoughts of relevance to this thesis. For an excellent overview and analysis of eleven models, see Miyake and Shah (1999).

Table 15. Models of WM.

	Model, Authors, & Select Citations	Summary of Model	Comments
1	Multi-Component WM Baddeley & Hitch (BH) (Baddeley & Hitch, 1974) (Baddeley, 2012)	Four components: - Central Executive - Visuospatial Sketchpad - Phonological Loop - Episodic Buffer (added later—see (Baddeley, Allen, & Hitch, 2011) The last three “slave” systems are considered “placeholders” with the possibility of others being added, especially to account for other sensory modalities.	Chiefly operations- focused.
2	Embedded Processes Model Cowan (Cowan, 1988) (Cowan, 1999)	An account that grew out of Baddeley & Hitch’s model, with five central features: ○ WM information comes from hierarchically arranges faculties: ▪ LTM ▪ Subset of LTM that is currently “activated”.	Chiefly operations- focused One problem with this model is that it doesn’t seem to account for data that never reaches

	Model, Authors, & Select Citations	Summary of Model	Comments
		<ul style="list-style-type: none"> ▪ Subset of activated memory that is “in the focus of attention and awareness”. ○ Focus of attention is capacity limited, activation is time limited. ○ Focus of attention is controlled conjointly by voluntary (central executive) and involuntary (orienting) systems. ○ There is habituation of orienting: stable stimuli activate memory but do not elicit awareness. ○ Awareness influences processing: it increases the resolution of perception and makes representations available for explicit recall. <p>For a diagrammatic representation of the model, see Cowan’s 1988, Figure 1, page 180 (reproduced following this table).</p>	LTM—perceptual or cognitive data that is used in WM then disappears forever. Of course, it wouldn’t be hard to tweak his model to accommodate this.
3	Mind and Brain Model Jonides (Jonides et al., 2008)	<p>An embedded processes model with significant detail and careful empirical grounding.</p> <p>As Baddeley (2012, p. 21) points out, this is a far more detailed approach than BH.</p>	Chiefly parts-focused, experimentally derived, especially neuroimaging.
4	Emergent Properties Model Postle	Alternative to a “standard model” of WM in which dedicated neurones in the PFC are the buffers of WM.	Chiefly parts-focused, experimentally derived.

	Model, Authors, & Select Citations	Summary of Model	Comments
	(Postle, 2006)	<p>Two chief principles (p.31):</p> <ul style="list-style-type: none"> ○ “the retention of information in working memory is associated with sustained activity in the same brain regions that are responsible for the representation of that information in non-working memory situations, such as perception, semantic memory, oculo- and skeletomotor control, and speech comprehension and production.” ○ “humans opportunistically, automatically, recruit as many mental codes as are afforded by a stimulus when representing that stimulus in working memory”. <p>Therefore, the PFC is not involved in the storage function of WM, but in the executive processing. He offers some specific possibilities, and notes that they are not unique to WM.</p>	
5	Facet Model Oberauer (Oberauer, 2009; Oberauer et al., 2003)	<p><i>2003 version:</i></p> <p>Two facets: function and domain:</p> <p>Three Functions:</p> <ul style="list-style-type: none"> ○ Simultaneous processing and storage ○ Supervision ○ Coordination <p>Two Domains:</p> <ul style="list-style-type: none"> ○ Visuo-spatial ○ Linguistic / Numeric <p>This gives six possible combinations.</p>	Chiefly operations-focused, experimentally derived.

	Model, Authors, & Select Citations	Summary of Model	Comments
		<p>Mainly aimed at providing a framework for useful experimental paradigms.</p> <p><i>2009 version:</i></p> <p>Three States of Representation:</p> <ul style="list-style-type: none"> ○ activated long-term memory ○ capacity limited region of direct access ○ focus of attention <p>Two Processes:</p> <ul style="list-style-type: none"> ○ Declarative ○ Procedural <p>Again, six possible combinations, with further subcomponents making the picture even more complex.</p>	
6	<p>Perceptual Memory Model</p> <p>Magnussen</p> <p>(Magnussen, 2009)</p>	<p>Five characteristics:</p> <ul style="list-style-type: none"> ○ Numerous, early level, highly perceptual stimulus-specific storage and processing units. ○ Inhibition occurs within units but not across units. ○ Each unit has its own independent resources. ○ Higher level processing brings together lower-order information. ○ Extraction of information is via a “cortical search” of lower level (V1) representations. <p>Further, low level representations may be encoded in LTM without ever attaining the level of semantic representation.</p>	<p>Chiefly operations-focused, experimentally derived.</p> <p>Highlights the fact that there is an implicit WM, and that the borders between implicit and explicit may be fuzzy.</p> <p>Is IWM something other than WM? But, no, it can be shown to have all the properties that define WM generally.</p>

	Model, Authors, & Select Citations	Summary of Model	Comments
			Highlights the extreme interrelatedness of cognitive processes generally.
7	Interactive Cognitive Subsystems Model (Barnard, 1999)	<p>A theoretical model that lends itself to computational modelling, based, as it is, on a multiprocessor architecture.</p> <p>The “core” theory assumes an architecture composed of nine subsystems. These all share an identical internal structure involving three basic resources (p. 299ff):</p> <ul style="list-style-type: none"> • processes that change the form in which information is represented; • an image record that preserves past input to the subsystem; and • a process that simply copies information into that record. 	Assessed positively in Logie & Cowan (2015, p. 317).
8	Computational Modelling See (Miyake & Shah, 1999, Chapters 5–11) ACT-R is mentioned by Jonides as a popular computational model.	A broad variety of more detailed accounts, involving computational modelling or simulation.	<p>Parts, operations and organisation-focused (comprehensive models).</p> <p>Often take a broader theoretical approach that can later be tested empirically by computer simulations</p>

	Model, Authors, & Select Citations	Summary of Model	Comments
	See (Lovett et al., 1999)		mimicking biological behaviour. Baddeley worries about the complexity of these models, although of Barnard's he says it "can also be mapped directly onto M-WM" (2012, p. 21).
9	Individual Difference Theories E.g., Engle et al. (Engle et al., 1999) (Engle & Kane, 2004)	<p>WM performance (retention <i>and</i> processing) seems to correlate with tasks like comprehension (i.e., better WM = greater comprehension), but <i>not</i> pure WM retention capacity. Why?</p> <p>On this basis, Engle et al distinguish WM from pure STM. They define WM succinctly as:</p> <p>"WM = STM (activated portion of LTM) + controlled attention" (1999, p.126).</p> <p>Exploring individual differences in performance can thus elucidate the nature of WM and its relationship to other cognitive processes.</p> <p>Rosenberg et al (2017, p. 299, figure 3) propose developing this approach using connectome-based modelling, in the same way they have modelled attention.</p>	<p>Chiefly operations-focused, experimentally derived.</p> <p>Highlights the interrelatedness of attention and LTM with WM.</p>

Appendix 6. Metaphors of WM: A Selection

Table 16. Metaphors of WM.

	Metaphors	Summary of Metaphors, with Citations	Comments
1	Storehouse / Box	WM is like a room or a box, into which are brought bits of data that are currently being manipulated, in order to achieve current goals. Having a limited amount of physical space in it, the room/box can only contain so many things at any one time (Fodor, 1983, pp. 8–9). The limited capacity may be measured by the number or items that can fit in the box, or more likely, by the precision with which items can be stored (see <i>Correspondence Models</i> above in Appendix 5), (Koriat, Goldsmith, & Pansky, 2000, pp. 482–485).	
2	Desktop Metaphor	WM is like a cluttered desktop, literal or on a computer, into which currently relevant documents are taken off the shelves and placed, consulted, and the information therein used to create new documents (my own invention).	

	Metaphors	Summary of Metaphors, with Citations	Comments
2	Telephone Exchange	WM is a pattern of one-to-one connections, driven by a central purposeful controller, in the same way that an old-fashioned telephone operator sat at a console and connected telephone customers to each other as appropriate. Useful connections are strengthened by the controller and thus more likely to be reused in future (Tolman, 1948, pp. 190–191).	
3	Map Control Room—Central Office / Global Workspace / Blackboard / Cauldron / Theatre	<p>This metaphor focuses on the manipulation aspect of WM: WM is like a map control room or a central office in which select, complex information is gathered and formed into a map, upon the basis of which, strategic decisions are taken (Tolman, 1948, p. 192); a workspace in which work is done (Baars, 1988); a blackboard on which calculations are performed (Oberauer, 2009, p. 86); a cauldron in which new memories are created from the right ingredients tossed in and stirred (Logie & Cowan, 2015, p. 320); or a theatre stage upon which the story is acted out (Baars, 1997a).</p> <p>“The focus of attention or episodic buffer might serve as a cauldron for the formation of new long-term memories” (Logie & Cowan, 2015, p. 320).</p>	<p>“In the working theatre, focal consciousness acts as a ‘bright spot’ on the stage, directed there by the selective ‘spotlight’ of attention. The bright spot is further surrounded by a ‘fringe,’ of vital but vaguely conscious events (Mangan, 1993). The entire stage of the theatre corresponds to ‘working memory’, the immediate memory system in which we talk to ourselves, visualize places and people, and plan actions” (Baars, 1997a, p. 292).</p>

	Metaphors	Summary of Metaphors, with Citations	Comments
4	Traffic Cop / Mental Energy / Resource Allocation	Here the focus is on the competition for cognitive processing resources or energy. WM is the “traffic cop”, directing the traffic of cognitive processing at a busy intersection of competing processes (Powledge, 1997, p. 333). What is being controlled and directed or allocated may be cognitive resources, or “mental energy” (see Miyake & Shah page 2, but they provide no citations).	
5	Juggler	A more entertaining version of the Traffic Cop Metaphor. WM juggles the competing processes in order to create something both beautiful and of practical use in achieving goals (Rossi, 1998, p. 118).	
6	Computer Hardware	Whereas LTM is roughly like the hard drive of a computer, WM is more like the RAM-processor assembly: a short term store of limited capacity (RAM) working together with a system for manipulating what is stored (CPU), that creates copies of data from LTM and manipulates them. Numerous computational WM models of varying degrees of complexity and detail have been proposed (see for example, Miyake & Shah, 1999).	

Appendix 7. Anatomy and Physiology of Foveal Vision

The Retina.

The fovea is a specialised region of the retina, temporal to the optic disc, that allows much higher resolution visual acuity than the rest of the retina.²⁹⁵ It does so by having a much higher density of photoreceptors, with the colour sensitive cones predominating over faint-light sensitive rods (the opposite pattern to the rest of the retina). Further, inner retinal neural elements are pushed aside, allowing more of the light to fall directly upon the receptors. Thus, retinal thickness is less at the fovea, creating the so-called “foveal pit,” which is gently sloped and subtends an angle on the retina of about 5°. This translates to the angle of the foveal visual field being 5° also. There is also an avascular zone subtending an angle of 2.5°, which also contributes to providing clear access for light to the receptor cells. At the bottom of the foveal pit lies the foveola, subtending an angle of 1.2°. The foveola contains only cones—no rods—at a density greater than anywhere else on the retina. Although the fovea represents 0.01% of the total area of the retina, approximately 8% of the striate visual cortex (primary visual cortex, or V1) is devoted to processing the information that derives from it.

Rods are more densely packed as one approaches the fovea, ranging from peripheral densities of about 50,000 rods/mm² to central densities of about 150,000 rods/mm² (p. 37). Cones, while distributed throughout the retina, have their highest densities in the fovea, where they can reach density of 225,000 cones/mm² at the centre. This high density, combined with the absence of neural and vascular elements, explains the much higher resolution possible in foveal vision. Further, rods are suited to

²⁹⁵ Unless otherwise stated, the information and the page references in this appendix are from Schwartz (2017, Chapters 2, 3).

scotopic light conditions (luminance of 10^{-6} – 10^{-4} candelas/m²) while cones are more suited to photopic light conditions (10^2 – 10^6 candelas/m²) (p. 27). Since almost all empirical paradigms exploring the relation between attention and consciousness involve photopic stimuli, it is cone vision that is of most interest, although the averted vision of astronomers viewing “faint fuzzies” (8.6.3) involves scotopic conditions.

What is the size of the foveal visual field? Clearly this is not a question with a straight answer. The fovea is not a discrete and uniform structure, but has gradations of cone density and neural and vascular paucity, increasing as one approaches the centre. The field of the foveola, which has maximum optimisation and resolution, is 1.2°, while that of the fovea proper is 5°. I present the relevant calculations for Ben’s hen below (Table 17 and Figure 8). In the main text, I will employ the foveal field size of 5°, although I do not think much would change even if we adopted the smaller field of the foveola for Ben’s hen, or if we accepted that the relevant resolution benefits extend further out than the fovea, albeit to gradually declining degree (see the discussion of Florence cases in 9.4.2). The foveal field covers a diameter of about 17cm at a distance of one metre, while the foveolar field covers a diameter of about 4.2cm at one-metre distance (see table below).

Table 17. Calculations of foveal and foveolar visual fields.

	Fovea	Foveola
Retinal angle	5°	1.2°
Size of field subtended at 1m from the eye	$r/1 = \tan 5^\circ$ $r = 0.087\text{m}$ $2 \times r = 0.17\text{m}$ Diameter of field = 17cm	$r/1 = \tan 1.2^\circ$ $r = 0.021\text{m}$ $2 \times r = 0.042\text{m}$ Diameter of field = 4.2cm

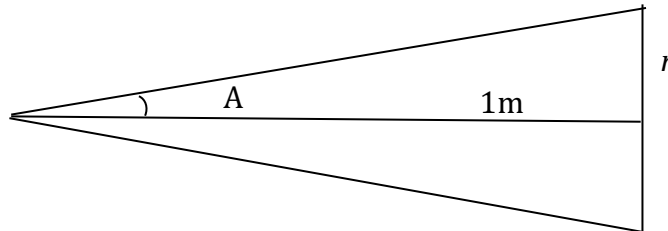


Figure 8 Visual field angle and diameter.

Resolution Acuity

The relative density of receptors in the foveola means that resolution is greatest there, gradually diminishing as one moves away from it. What the eye can resolve depends on many complex interacting factors (Schwarz, Chapter 7). Factors that impact on resolution acuity include the degree of contrast in the stimulus (p.158, Figures 7-6 and 7-7) and of course, focus (p. 156, Figure 7-5). For a relatively high contrast speckled hen in good focus, being viewed by a young healthy adult, the limit of resolution acuity is approximately 60 cycles/degree (p. 157). That is, young and healthy Ben can resolve a sine-wave grating with 60 dark-light bands per degree of visual field. If we assume each dark-light band to equate to a single speckle, then the limit of Ben's resolution acuity is to resolve 60 speckles across each degree of his foveolar field. Given that field subtends an angle of 1.2° , Ben could resolve up to 72 speckles laid out in a straight line across his foveolar field. A speckled hen with 100 speckles laid out in a 10×10 grid would be comfortably within Ben's acuity limits to resolve.

In summary, a speckled hen fitting within a 17cm diameter circle at a distance of one metre would fit completely within Ben's foveal field, easily allowing him to resolve every speckle with great clarity. A hen fitting inside a 4.2cm diameter circle at one

metre distance would fit completely within Ben's foveolar field, also easily allowing him to resolve every speckle clearly.

Post-Retina.

Axons in the optic nerve, bringing information from the retina, mostly terminate in the Lateral Geniculate Nucleus (LGN) of the thalamus. From there, cells communicate mostly with V1, and from there, to higher visual processing regions (collectively known as the extrastriate cortex) and to many other regions of the brain. there are also significant tracts connecting these regions in the opposite directions. The retinocortical projection (retina to LGN to V1) has three distinct parallel pathways: *parvo* (spatial and red colour information); *magno* (fast movement); and *konio* (blue-yellow colour).

Signals from the retina undergo a kind of convergence. The roughly 100,000,000 photoreceptors (120,000,000 rods and 6,000,000 cones, p. 37) converge upon 1,000,000 ganglion cells. However, in the cortex, signals undergo significant divergence, spreading out from V1 first to the extrastriate cortex that are specialised for analysing attributes such as motion and colour, and then "to higher centres, which combine visual information with memory and other senses. Higher visual centres, in turn, send information back to the striate cortex (V1) via reciprocal projections" (p. 20). All of this falls under my definition of Liberal Attention, of course.

Studies have shown that individual cells in V4 show selective levels of activation depending on which object a monkey is paying attention to, even if the eye and the stimuli are in exactly the same position. On the other hand, cells in V1 do not show such selective activation due to object attention.²⁹⁶ This suggests strongly that visual attention is mediated by the extrastriate cortex, and not by V1. "This suggests that there may be a filtering of visual information from striate cortex to V4 that is at least

²⁹⁶ Although V1 cell activation may be modulated by *spatial* attention (Alex Holcombe, personal communication).

partially dependent on the stimulus to which the animal is attending" (p. 305). Also, in footnote 2 on the same page: "The receptive fields of neurons in higher visual centres tend to be larger than those in lower centres, making it possible to study them with large, complex stimulus arrays." Does this not argue for the richness of FOC?

Bibliography

- Ackerman, P. L., Beier, M. E., & Boyle, M. O. (2005). Working Memory and Intelligence: The Same or Different Constructs? *Psychological Bulletin*, 131(1), 30–60.
- Adamson, P., & Ganeri, J. (2017). *The Truth Shall Set You Free: Nyāya on the Mind*. History of Philosophy in India Podcast. Retrieved from <https://historyofphilosophy.net/nyaya-mind>
- Agrillo, C. (2011). Near-Death Experience: Out-of-Body and Out-of-Brain? *Review of General Psychology*, 15(1), 1–10. <http://doi.org/10.1037/a0021992>
- Ahmadi, M. R., Gilakjani, A. P., & Ahmadi, S. M. (2011). The Relationship between Attention and Consciousness. *Journal of Language Teaching & Research*, 2(6), 1366–1373. <http://doi.org/10.4304/jltr.2.6.1366-1373>
- Aiello, L. C., & Wheeler, P. (1995). The expensive-tissue hypothesis: the brain and the digestive system in human and primate evolution. *Current Anthropology*, 36(2), 199–221.
- Alcaraz, A. (2018). Looking for the minimal phenomenal experience: the case of lucid dreamless sleep. Retrieved October 20, 2018, from https://www.academia.edu/37140556/The_case_of_lucid_dreamless_sleep
- Allport, A. (1980). Attention and Performance. In G. Claxton (Ed.), *Cognitive Psychology: New Directions* (pp. 112–153). Routledge.
- Allport, A. (1987). Selection for Action. In H. Heuer & H. Sanders (Eds.), *Perspectives on Perception and Action* (pp. 395–419). Lawrence Erlbaum Associates Inc.
- Allport, A. (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In D. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 183–218). MIT Press.
- Alvarez-Silva, S., Alvarez-Rodriguez, J., & Cavanna, A. E. (2012). Epileptic aura and qualitative alterations of consciousness in focal seizures: A neuropsychiatric approach. *Epilepsy & Behavior*, 23(4), 512–513. <http://doi.org/https://doi.org/10.1016/j.yebeh.2011.11.031>
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15(3), 122–131. <http://doi.org/10.1016/j.tics.2011.01.003>
- Analayo, B. (2012). *Excursions into the Thought-World of the Pali Discourses*. Pariyatti Publishing.
- Anderson, B. L., & Winawer, J. (2005). Image segmentation and lightness perception. *Nature*, 434(3 March), 79–83. <http://doi.org/10.1038/nature03343>.

- Anderson, M. L. (2010). Neural reuse: a fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(4), 245–266.
- Anderson, M. L., Kinnison, J., & Pessoa, L. (2013). Describing functional diversity of brain regions and brain networks. *NeuroImage*, 73, 50–58.
<http://doi.org/10.1016/j.neuroimage.2013.01.071>
- Andrews, K. (2016). Animal Cognition. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/sum2016/entries/cognition-animal/>
- Andrillon, T., & Kouider, S. (2016). Implicit memory for words heard during sleep. *Neuroscience of Consciousness*, (1), 1–12. <http://doi.org/10.1093/nc/niw014>
- Armstrong, D. M. (1995). Perception-consciousness and action-consciousness? *Behavioral and Brain Sciences*, 18(2), 247–248.
- Aru, J., & Bachmann, T. (2017). Expectation creates something out of nothing: The role of attention in iconic memory reconsidered. *Consciousness and Cognition*, 53, 203–210.
- Arvidson, P. S. (2003). A lexicon of attention: from cognitive science to phenomenology. *Phenomenology and the Cognitive Sciences*, 2(2), 99–132.
<http://doi.org/10.1023/A:1024895827774>
- Augustine. (1950). *Against the Academicians*. (J. Quasten & J. C. Plumpe, Eds., J. J. O'Meara, Trans.). Newman Press.
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual Working Memory Represents a Fixed Number of Items Regardless of Complexity. *Psychological Science*, 18(7), 622–628.
<http://doi.org/10.1111/j.1467-9280.2007.01949.x>
- Awh, E., Vogel, E., & Oh, S.-H. (2006). Interactions between attention and working memory. *Neuroscience*, 139(1), 201–208.
<http://doi.org/https://doi.org/10.1016/j.neuroscience.2005.08.023>
- Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Baars, B. J. (1997a). In the Theatre of Consciousness. Global Workspace Theory, A Rigorous Scientific Theory of Consciousness. *Journal of Consciousness Studies*, 4(4), 292–309.
- Baars, B. J. (1997b). Some Essential Differences between Consciousness and Attention, Perception, and Working Memory. *Consciousness and Cognition*, 6(2–3), 363–371.
<http://doi.org/10.1006/ccog.1997.0307>
- Baars, B. J. (2013). A scientific approach to silent consciousness. *Frontiers in Psychology*, 4(678), 1–3. <http://doi.org/10.3389/fpsyg.2013.00678>
- Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences*, 7(4), 166–172. [http://doi.org/10.1016/S1364-6613\(03\)00056-1](http://doi.org/10.1016/S1364-6613(03)00056-1)
- Baddeley, A. D. (1993a). Working memory and conscious awareness. In A. F. Collins, S. E. Gathercole, M. A. Conway, & P. E. Morris (Eds.), *Theories of Memory* (pp. 11–28). Lawrence Erlbaum Associates.
- Baddeley, A. D. (1993b). Working Memory or Working Attention? In A. D. Baddeley & L. Weiskrantz (Eds.), *Attention: Selection, awareness and control. A tribute to Donald Broadbent* (pp. 152–170). Oxford; New York: Clarendon.
- Baddeley, A. D. (1996). Exploring the Central Executive. *The Quarterly Journal of Experimental Psychology Section A*, 49(1), 5–28.
- Baddeley, A. D. (2000). The episodic buffer: a new component of working memory? *Trends*

- in Cognitive Sciences*, 4(11), 417–423. [http://doi.org/10.1016/S1364-6613\(00\)01538-2](http://doi.org/10.1016/S1364-6613(00)01538-2)
- Baddeley, A. D. (2003). Working memory: looking back and looking forward. *Nature Reviews Neuroscience*, 4(10), 829–839. <http://doi.org/10.1038/nrn1201>
- Baddeley, A. D. (2007). *Working Memory, Thought, and Action*. Oxford University Press.
- Baddeley, A. D. (2010). Working memory. *Current Biology*, 20(4), R136–R140. <http://doi.org/10.1016/j.cub.2009.12.014>
- Baddeley, A. D. (2012). Working Memory: Theories, Models, and Controversies. *Annual Review of Psychology*, 63(1), 1–29.
- Baddeley, A. D., Allen, R. J., & Hitch, G. J. (2011). Binding in visual working memory: The role of the episodic buffer. *Neuropsychologia*, 49(6), 1393–1400. <http://doi.org/10.1016/j.neuropsychologia.2010.12.042>
- Baddeley, A. D., & Andrade, J. (2000). Working Memory and the Vividness of Imagery. *Journal of Experimental Psychology: General*, 129(1), 126–145. <http://doi.org/10.1037/0096-3445.129.1.126>
- Baddeley, A. D., & Hitch, G. J. (1974). Working Memory. In G. Bower (Ed.), *Recent advances in learning and motivation, Volume 8* (pp. 47–90). Academic Press.
- Baillet, S. (2017). Magnetoencephalography for brain electrophysiology and imaging. *Nature Neuroscience*, 20(3), 327–339. <http://doi.org/10.1038/nn.4504>
- Barak, O., & Tsodyks, M. (2014). Working models of working memory. *Current Opinion in Neurobiology*, 25, 20–24. <http://doi.org/10.1016/j.conb.2013.10.008>
- Barbot, A., & Carrasco, M. (2017). Attention Modifies Spatial Resolution According to Task Demands. *Psychological Science*, 28(3), 285–296. <http://doi.org/10.1177/0956797616679634>
- Barinaga, M. (2002). How the Brain's Clock Gets Daily Enlightenment. *Science*, 295(5557), 955–957.
- Barnard, P. J. (1999). Interacting Cognitive Subsystems: Modelling Working Memory Phenomena Within a Multiprocessor Architecture. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 298–339). Cambridge University Press.
- Barrett, D. (2014). Consciousness, Attention, and Working Memory. *Journal of Consciousness Studies*, 21(9–10), 7–29.
- Barron, A. B., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, 113(18), 4900–4908. <http://doi.org/10.1073/pnas.1520084113>
- Bayliss, D. M., Jarrold, C., Gunn, D. M., & Baddeley, A. D. (2003). The Complexities of Complex Span: Explaining Individual Differences in Working Memory in Children and Adults. *Journal of Experimental Psychology: General*, 132(1), 71–92. <http://doi.org/10.1037/0096-3445.132.1.71>
- Bayne, T. (2005). Divided brains and unified phenomenology: a review essay on Michael Tye's *Consciousness and Persons*. *Philosophical Psychology*, 18(4), 495–512. <http://doi.org/10.1080/09515080500229993>
- Bayne, T. (2010). *The Unity of Consciousness*. Oxford University Press.
- Bayne, T. (2013). Agency as a Marker of Consciousness. In T. Vierkant, J. Kiverstein, & A. Clark (Eds.), *Decomposing the Will* (pp. 160–180). Oxford University Press.

<http://doi.org/10.1093/ACPROF:OSO/9780199746996.003.0009>

- Bayne, T., & Chalmers, D. J. (2003). What is the Unity of Consciousness? In A. Cleeremans (Ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. (pp. 23–58). Oxford University Press.
- Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences*, 20(6), 405–413. <http://doi.org/10.1016/j.tics.2016.03.009>
- Bayne, T., & Montague, M. (Eds.). (2011). *Cognitive Phenomenology*. Oxford University Press.
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49(6), 1622–1631. <http://doi.org/https://doi.org/10.1016/j.neuropsychologia.2010.12.023>
- Beaman, C. P. (2010). Working memory and working attention: What could possibly evolve? *Current Anthropology*, 51(S1), S27–S38. <http://doi.org/10.1086/650297>
- Bechtel, W. (2001). The Compatibility of Complex Systems and Reduction: A Case Analysis of Memory Research. *Minds and Machines*, 11(4), 483–502. <http://doi.org/10.1023/A:1011803931581>
- Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. Routledge.
- Bechtel, W., Abrahamsen, A., & Graham, G. (2001). Cognitive Science: History. In N. Smelser & P. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences* (pp. 2154–2158). Pergamon.
- Bello, P. F., & Bridewell, W. (2017). There Is No Agency Without Attention. *AI Magazine*, 38(4), 27–33.
- Bender, S. (2016). Reflection and Rationality in Leibniz. In J. Kaukua & T. Ekenberg (Eds.), *Subjectivity and Selfhood in Medieval and Early Modern Philosophy* (pp. 263–275). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-26914-6_16
- Berg, A. T., Berkovic, S. F., Brodie, M. J., Buchhalter, J., Cross, J. H., Van Emde Boas, W., ... Scheffer, I. E. (2010). Revised terminology and concepts for organization of seizures and epilepsies: Report of the ILAE Commission on Classification and Terminology, 2005–2009. *Epilepsia*, 51(4), 676–685. <http://doi.org/10.1111/j.1528-1167.2010.02522.x>
- Bering, J. M., & Bjorklund, D. F. (2007). The Serpent's Gift: Evolutionary Psychology and Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness* (pp. 597–629). Cambridge University Press.
- Berrios, G. E. (1981). Stupor: a conceptual history. *Psychological Medicine*, 11(4), 677–688.
- Bickle, J. (2016). Multiple Realizability. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2016/entries/multiple-realizability/>
- Binder, M. D., Hirokawa, N., & Windhorst, U. (Eds.). (2009). Stretch Receptor. In *Encyclopedia of Neuroscience* (p. 3865). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-29678-2_5689
- Bitter, D. (2010). *The Dissociation of Consciousness*. Central European University. Retrieved from http://www.etd.ceu.hu/2010/bitter_david.pdf
- Blackmore, S. (1991). Near-Death Experiences: In or out of the body? *Skeptical Inquirer*, 16, 34–45.

- Bleicher, A. (2012). Edges of Perception: Unusual cases reveal that the famous “five senses” are not as distinct as once thought. *Scientific American Mind*, 23(1), 46–53.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227–247. <http://doi.org/10.1017/S0140525X00038188>
- Block, N. (2001). Paradox and cross purposes in recent work on consciousness. *Cognition*, 79(1), 197–219. [http://doi.org/10.1016/S0010-0277\(00\)00129-3](http://doi.org/10.1016/S0010-0277(00)00129-3)
- Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9(2), 46–52. <http://doi.org/10.1016/j.tics.2004.12.006>
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5–6), 481–99; discussion 499–548. <http://doi.org/10.1017/S0140525X07002786>
- Block, N. (2011a). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15(12), 567–575. <http://doi.org/10.1016/j.tics.2011.11.001>
- Block, N. (2011b). The higher order approach to consciousness is defunct. *Analysis*, 71(3), 419–431.
- Block, N. (2013a). Seeing and Windows of Integration. *Thought: A Journal of Philosophy*, 2(1), 29–39. <http://doi.org/10.1002/tht.62>
- Block, N. (2013b). The Grain of Vision and the Grain of Attention. *Thought*, 1, 170–184. <http://doi.org/10.1002/tht.28>
- Block, N. (2014a). Rich conscious perception outside focal attention. *Trends in Cognitive Sciences*, 18(9), 445–447. <http://doi.org/10.1016/j.tics.2014.05.007>
- Block, N. (2014b). The Defective Armchair: A Reply to Tye. *Thought: A Journal of Philosophy*, 3(2), 159–165. <http://doi.org/10.1002/tht3.124>
- Block, N. (2018). If perception is probabilistic, why does it not seem probabilistic? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170341.
- Block, N., Carmel, D., Fleming, S. M., Kentridge, R. W., Koch, C., Lamme, V. A., ... Rosenthal, D. M. (2014). Consciousness science: real progress and lingering misconceptions. *Trends in Cognitive Sciences*, 18(11), 556–557. <http://doi.org/10.1016/j.tics.2014.09.004>
- Bloom, P. (2017). Empathy and Its Discontents. *Trends in Cognitive Sciences*, 21(1), 24–31. <http://doi.org/10.1016/j.tics.2016.11.004>
- Blume, W. T., Lüders, H. O., Mizrahi, E., Tassinari, C., Van Emde Boas, W., & Engel Jr., J. (2001). Glossary of descriptive terminology for ictal semiology: Report of the ILAE Task Force on classification and terminology. *Epilepsia*, 42(9), 1212–1218. <http://doi.org/10.1046/j.1528-1157.2001.22001.x>
- Bogen, J. E. (2007). The Thalamic Intralaminar Nuclei and the Property of Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 775–807). Cambridge University Press.
- Bourget, D. (2017). Why Are Some Phenomenal Experiences ‘Vivid’ and Others ‘Faint’? Representationalism, Imagery, and Cognitive Phenomenology. *Australasian Journal of Philosophy*, 95(4), 673–687. <http://doi.org/10.1080/00048402.2017.1278612>
- Bourget, D., & Chalmers, D. J. (2014). What do philosophers believe? *Philosophical Studies*, 170(3), 465–500. <http://doi.org/10.1007/s11098-013-0259-7>
- Bradley, F. (1886). Is there any special activity of attention? *Mind*, 11(43), 305–323.
- Brady, T. F., & Alvarez, G. A. (2015). No Evidence for a Fixed Object Limit in Working

- Memory: Spatial Ensemble Representations Inflate Estimates of Working Memory Capacity for Complex Objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(3), 921–929.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487–502. <http://doi.org/10.1037/a0016797>
- Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *Proceedings of the National Academy of Sciences of the United States of America*, 113(27), 7459–7464. <http://doi.org/10.1073/pnas.1520027113>
- Brainerd, C. J., & Reyna, V. F. (2005). *The science of false memory*. Oxford University Press.
- Braun, J., & Sagi, D. (1990). Vision outside the focus of attention. *Perception & Psychophysics*, 48(1), 45–58. <http://doi.org/10.3758/BF03205010>
- Brenner, R. P. (2005). The Interpretation of the EEG in Stupor and Coma. *The Neurologist*, 11(5), 271–284.
- Broadbent, D. (1958). *Perception and Communication*. Pergamon Press.
- Broadbent, D. (1971). *Decision and stress*. Academic Press.
- Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We See More Than We Can Report: “Cost Free” Color Phenomenality Outside Focal Attention. *Psychological Science*, 25(7), 1394–1403. <http://doi.org/10.1177/0956797614532656>
- Bronfman, Z. Z., Jacobson, H., & Usher, M. (2018). Impoverished or rich consciousness outside attentional focus: Recent data tip the balance for Overflow. *Mind & Language*. <http://doi.org/10.1111/mila.12217>
- Brook, A., & Raymont, P. (2014). The Unity of Consciousness. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Spring 201). Stanford University. Retrieved from <http://plato.stanford.edu/archives/spr2014/entries/consciousness-unity/>
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain’s default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124, 1–38.
- Burgess, N., & Hitch, G. J. (2005). Computational models of working memory: putting long-term memory into context. *Trends in Cognitive Sciences*, 9(11), 535–541. <http://doi.org/10.1016/j.tics.2005.09.011>
- Cambridge Declaration on Consciousness*. (2012). Cambridge, UK. Retrieved from <http://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>
- Campana, F., & Tallon-Baudry, C. (2013). Anchoring visual subjective experience in a neural model: The coarse vividness hypothesis. *Neuropsychologia*, 51(6), 1050–1060. <http://doi.org/10.1016/j.neuropsychologia.2013.02.021>
- Campbell, J. (2011). Visual Attention and the Epistemic Role of Consciousness. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 323–341). Oxford University Press.
- Campos, A., & Pérez, M. J. (1988). Visual Elaboration Scale as a Measure of Imagery. *Perceptual and Motor Skills*, 66(2), 411–414. <http://doi.org/10.2466/pms.1988.66.2.411>
- Carhart-Harris, R. L., Erritzoe, D., Williams, T., Stone, J. M., Reed, L. J., Colasanti, A., ... Nutt, D. J. (2012). Neural correlates of the psychedelic state as determined by fMRI studies with psilocybin. *Proceedings of the National Academy of Sciences*, 109(6), 2138–2143.

- Carp, J. (2012). On the Plurality of (Methodological) Worlds: Estimating the Analytic Flexibility of fMRI Experiments. *Frontiers in Neuroscience*, 6(149). <http://doi.org/10.3389/fnins.2012.00149>
- Carr, N. G. (2010). *The Shallows : How the Internet is Changing the Way We Think, Read and Remember*. Atlantic Books.
- Carrasco, M., & Barbot, A. (2019). Spatial attention alters visual appearance. *Current Opinion in Psychology*, 29, 56–64. <http://doi.org/10.1016/j.copsyc.2018.10.010>
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, 7(3), 308–313. <http://doi.org/10.1038/nn1194>
- Carruthers, P. (2016). Higher-Order Theories of Consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2016). Metaphysics Research Lab, Stanford University.
- Carruthers, P. (2017). Block's Overflow Argument. *Pacific Philosophical Quarterly*, 98(S1), 65–70. <http://doi.org/10.1111/papq.12152>
- Cartwright, R. (2004). Sleepwalking Violence: A Sleep Disorder, a Legal Dilemma, and a Psychological Challenge. *American Journal of Psychiatry*, 161(7), 1149–1158. <http://doi.org/10.1176/appi.ajp.161.7.1149>
- Chadha, M. (2015). Meditation and unity of consciousness: a perspective from Buddhist epistemology. *Phenomenology and the Cognitive Sciences*, 14, 111–127. <http://doi.org/10.1007/s11097-013-9316-0>
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. J. (1997). Availability: The Cognitive Basis of Experience? In N. Block, O. Flanagan, & G. Guzeldere (Eds.), *The Nature of Consciousness: Philosophical Debates* (pp. 421–424). MIT Press.
- Chalmers, D. J. (2000). What Is a Neural Correlate of Consciousness? In T. Metzinger (Ed.), *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (pp. 17–39). MIT Press. <http://doi.org/10.1017/CBO9781107415324.004>
- Chalmers, D. J. (2010). *The Character of Consciousness*. Oxford University Press.
- Chalmers, D. J. (2011). Verbal Disputes. *Philosophical Review*, 120(4), 515–566. <http://doi.org/10.1215/00318108-1334478>
- Chappell, R. Y., & Yetter-Chappell, H. (2016). Virtue and Salience. *Australasian Journal of Philosophy*, 94(3), 449–463. <http://doi.org/10.1080/00048402.2015.1115530>
- Cherniak, C., Mokhtarzada, Z., Rodriguez-Esteban, R., & Changizi, K. (2004). Global optimization of cerebral cortex layout. *Proceedings of the National Academy of Sciences*, 101(4), 1081–1086. <http://doi.org/10.1073/pnas.0305212101>
- Chica, A. B., Botta, F., Lupiáñez, J., & Bartolomeo, P. (2012). Spatial attention and conscious perception: Interactions and dissociations between and within endogenous and exogenous processes. *Neuropsychologia*, 50, 621–629. <http://doi.org/10.1016/j.neuropsychologia.2011.12.020>
- Chou, W.-L., & Yeh, S.-L. (2012). Object-based attention occurs regardless of object awareness. *Psychonomic Bulletin & Review*, 19(2), 225–231.
- Christensen, W., Sutton, J., & McIlwain, D. J. F. (2016). Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience. *Mind & Language*, 31(1), 37–66. <http://doi.org/10.1111/mila.12094>

- Christoff, K., Mills, C., Andrews-Hanna, J. R., Irving, Z. C., Thompson, E., Fox, K. C. R., & Kam, J. W. Y. (2018). Mind-Wandering as a Scientific Concept: Cutting through the Definitional Haze. *Trends in Cognitive Sciences*, 22(11), 957–959.
<http://doi.org/10.1016/j.tics.2018.07.004>
- Christophel, T. B., Klink, P. C., Spitzer, B., Roelfsema, P. R., & Haynes, J.-D. (2017). The Distributed Nature of Working Memory. *Trends in Cognitive Sciences*, 21(2), 111–124.
<http://doi.org/10.1016/j.tics.2016.12.007>
- Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, 49(6), 1407–1409.
<http://doi.org/10.1016/j.neuropsychologia.2011.01.029>
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A Taxonomy of External and Internal Attention. *Annual Review of Psychology*, 62(1), 73–101.
<http://doi.org/10.1146/annurev.psych.093008.100427>
- Churchland, P. S. (1983). Consciousness: The transmutation of a concept. *Pacific Philosophical Quarterly*, 64(1), 80–95.
- Ciaramelli, E., Grady, C. L., & Moscovitch, M. (2008). Top-down and bottom-up attention to memory: A hypothesis (AtoM) on the role of the posterior parietal cortex in memory retrieval. *Neuropsychologia*, 46(7), 1828–1851.
<http://doi.org/10.1016/j.neuropsychologia.2008.03.022>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action and the Embodied Mind*. Oxford University Press.
- Clark, R. N. (1990). *Visual Astronomy of the Deep Sky*. Cambridge University Press and Sky Publishing Corporation.
- Cohen, M. A. (2019). What Is the True Capacity of Visual Cognition? *Trends in Cognitive Sciences*, 23(2), 83–86. <http://doi.org/10.1016/j.tics.2018.12.002>
- Cohen, M. A., Alvarez, G. A., & Nakayama, K. (2011). Natural-Scene Perception Requires Attention. *Psychological Science*, 22(9), 1165–1172.
<http://doi.org/10.1177/0956797611419168>
- 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>
- Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15(8), 358–364.
<http://doi.org/10.1016/j.tics.2011.06.008>
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016a). Ensemble Perception, Summary Statistics, and Perceptual Awareness: A Response. *Trends in Cognitive Sciences*, 20(9), 643–644. <http://doi.org/10.1016/j.tics.2016.06.007>
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016b). What is the Bandwidth of Perceptual Experience? *Trends in Cognitive Sciences*, 20(5), 324–335.
<http://doi.org/10.1016/j.tics.2016.03.006>
- Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics*, 27(3), 183–228. <http://doi.org/10.3758/BF03204258>
- Combs, A., Krippner, S., & Taylor, E. (2010). Is there Awareness Outside Attention? *Journal of Consciousness Studies*, 17(11–12), 100–115.
- Constantinidis, C., Funahashi, S., Lee, D., Murray, J. D., Qi, X.-L., Wang, M., & Arnsten, A. F. T. (2018). Persistent Spiking Activity Underlies Working Memory. *The Journal of*

- Neuroscience*, 38(32), 7020–7028. <http://doi.org/10.1523/JNEUROSCI.2486-17.2018>
- Corballis, M. C. (2007). The Evolution of Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness* (pp. 571–595). Cambridge University Press.
- Corbetta, M., & Shulman, G. (2002). Control of Goal-Directed and Stimulus-Driven Attention in the Brain. *Nature Reviews*, 3(3), 215–229.
- Corkin, S. (2002). What's new with the amnesic patient H.M.? *Nature Reviews Neuroscience*, 3, 153. <http://doi.org/10.04.14/nrn726>
- Cosmelli, D., Lachaux, J.-P., & Thompson, E. (2007). Neurodynamical Approaches to Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 731–772). Cambridge University Press.
- Cowan, N. (1988). Evolving Conceptions of Memory Storage, Selective Attention, and Their Mutual Constraints Within the Human Information-Processing System. *Psychological Bulletin*, 104(2), 163–191.
- Cowan, N. (1995). *Attention and memory: An integrated framework*. Oxford University Press.
- Cowan, N. (1999). An Embedded-Processes Model of Working Memory. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 62–101). Cambridge University Press.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87–114.
- Cowan, N. (2005). *Working Memory Capacity*. Psychology Press.
- Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? *Progress in Brain Research*, 169, 323–338.
- Cowan, N. (2017). The many faces of working memory and short-term storage. *Psychonomic Bulletin & Review*, 24(4), 1158–1170. <http://doi.org/10.3758/s13423-016-1191-6>
- Cowan, N., Elliott, E. M., Scott Saults, J., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. A. (2005). On the Capacity of Attention: Its Estimation and Its Role in Working Memory and Cognitive Aptitudes. *Cognitive Psychology*, 51(1), 42–100. <http://doi.org/10.1016/j.cogpsych.2004.12.001>
- Cowan, N., & Morey, C. C. (2006). Visual working memory depends on attentional filtering. *Trends in Cognitive Sciences*, 10(4), 139–141. <http://doi.org/10.1016/j.tics.2006.02.001>
- Crick, F. C., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2, 263–275. Retrieved from <http://papers.klab.caltech.edu/22/1/148.pdf>
- Crick, F. C., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, 6(2), 119–126. <http://doi.org/10.1002/9780470751466.ch45>
- Crowder, R. G. (1982). The demise of short-term memory. *Acta Psychologica*, 50(3), 291–323. [http://doi.org/10.1016/0001-6918\(82\)90044-0](http://doi.org/10.1016/0001-6918(82)90044-0)
- Cudworth, R. (1678). *The True Intellectual System of The Universe*.
- Cunningham, C. C., Chen, W. P., Shorten, A. T., McClurkin, M. T., Choezom, T. T., Schmidt, C. T., ... Blumenfeld, H. T. (2014). Impaired consciousness in partial seizures is bimodally distributed. *Neurology*, 82(19), 1736–1744. <http://doi.org/10.1212/WNL.0000000000000404>
- Dainton, B. (2000). *Stream of Consciousness: unity and continuity in conscious experience*.

Routledge.

- Dainton, B. (2008). *The Phenomenal Self*. Oxford University Press.
<http://doi.org/10.1093/acprof:oso/9780199288847.001.0001>
- Dalrymple, K. A., Barton, J. J., & Kingstone, A. (2013). A world unglued: simultanagnosia as a spatial restriction of attention. *Frontiers in Human Neuroscience*, 7(145), 1–10.
<http://doi.org/10.3389/fnhum.2013.00145>
- Danziger, K. (2008). *Marking the mind: A history of memory*. Cambridge University Press.
- De-Miguel, F. F., & Fuxe, K. (2012). Extrasynaptic neurotransmission as a way of modulating neuronal functions. *Frontiers in Physiology*, 3(16), 1–2.
<http://doi.org/10.3389/fphys.2012.00016>
- De-Wit, L., Alexander, D., Ekroll, V., & Wagemans, J. (2016). Is neuroimaging measuring information in the brain? *Psychonomic Bulletin & Review*, 1–14.
<http://doi.org/10.3758/s13423-016-1002-0>
- De Brigard, F. (2010). Consciousness, Attention and Commonsense. *Journal of Consciousness Studies*, 17(9–10), 189–201.
- De Brigard, F., & Prinz, J. (2010). Attention and consciousness. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 51–59. <http://doi.org/10.1002/wcs.27>
- De Deyne, C., Struys, M., Decruyenaere, J., Creupelandt, J., Hoste, E., & Colardyn, F. (1998). Use of continuous bispectral EEG monitoring to assess depth of sedation in ICU patients. *Intensive Care Medicine*, 24(12), 1294–1298.
<http://doi.org/10.1007/s001340050765>
- de Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, 18(3), 569–577.
<http://doi.org/https://doi.org/10.1016/j.concog.2009.03.002>
- de Graaf, T. A., Hsieh, P.-J., & Sack, A. T. (2012). The ‘correlates’ in neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews*, 36(1), 191–197.
<http://doi.org/10.1016/j.neubiorev.2011.05.012>
- De Loof, E., Poppe, L., Cleeremans, A., Gevers, W., & Van Opstal, F. (2015). Different effects of executive and visuospatial working memory on visual consciousness. *Attention, Perception, & Psychophysics*, 77(8), 2523–2528. <http://doi.org/10.3758/s13414-015-1000-8>
- De Ridder, D., & Vanneste, S. (2013). The Artful Mind: Sexual selection and an evolutionary neurobiological approach to aesthetic appreciation. *Perspectives in Biology and Medicine*, 56(3), 327–340. <http://doi.org/10.1353/pbm.2013.0029>
- De Sousa, R. (2002). Fringe consciousness and the multifariousness of emotions. *Psyche*, 8(14).
- de Vignemont, F. (2014). Multimodal Unity and Multimodal Binding. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 125–150). MIT Press.
- 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), 204–11. <http://doi.org/10.1016/j.tics.2006.03.007>
- Dehaene, S., Charles, L., King, J. R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76–84.
<http://doi.org/10.1016/j.conb.2013.12.005>
- Dehaene, S., Lau, H., & Kouider, S. (2017). What is consciousness, and could machines have it? *Science*, 358, 486–492. <http://doi.org/10.1126/science.aan8871>

- Dennett, D. C. (1991). *Consciousness Explained*. Penguin Books.
- Dennett, D. C. (1995). Animal consciousness and why it matters. *Social Research*, 62, 691–710.
- Dennett, D. C. (2003). Who's On First? Heterophenomenology Explained. *Journal of Consciousness Studies*, 10(9–10), 19–30. <http://doi.org/10.1630/0326785041834748>
- Dennett, D. C. (2006). Cognitive Wheels: The Frame Problem of AI. In J. L. Bermudez (Ed.), *Philosophy of Psychology: Contemporary Readings* (pp. 433–454). Routledge.
- Dennett, D. C. (2007). Heterophenomenology reconsidered. *Phenomenology and the Cognitive Sciences*, 6(1–2), 247–270. <http://doi.org/10.1007/s11097-006-9044-9>
- Descartes, R. (1637). *Discourse on the Method*.
- Descartes, R. (1988). Replies to Objections. In J. Cottingham, R. Stoothoff, D. Murdoch, & A. Kenny (Eds.), *The Philosophical Writings of Descartes, volume II*. Cambridge University Press.
- Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, 18, 193–222. <http://doi.org/10.1146/annurev.ne.18.030195.001205>
- Desimone, R., & Moran, J. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782–784.
- Deutsch, R., Gawronski, B., & Strack, F. (2006). At the boundaries of automaticity: Negation as reflective operation. *Journal of Personality and Social Psychology*, 91(3), 385–405. <http://doi.org/10.1037/0022-3514.91.3.385>
- Devinsky, O., & Luciano, D. (1991). Psychic phenomena in partial seizures. In *Seminars in Neurology* (Vol. 11, pp. 100–109).
- Dijksterhuis, A., & Aarts, H. (2010). Goals, Attention, and (Un)Consciousness. *Annual Review of Psychology*, 61(1), 467–490. <http://doi.org/10.1146/annurev.psych.093008.100445>
- Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science*, 1, 95–109.
- Doyen, S., Klein, O., Simons, D. J., & Cleeremans, A. (2014). On the other side of the mirror: Priming in cognitive and social psychology. *Social Cognition*, 32(Supplement), 12–32.
- Dretske, F. (1993). Conscious Experience. *Mind*, 102(406), 263–283. Retrieved from <http://dx.doi.org/10.1093/mind/102.406.263>
- Dreyfus, G., & Thompson, E. (2007). Asian Perspectives: Indian Theories of Mind. In P. D. Zelazo, M. Moskovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 89–116). Cambridge University Press.
- Droege, P. (2009). Now or never: how consciousness represents time. *Consciousness and Cognition*, 18(1), 78–90. <http://doi.org/10.1016/j.concog.2008.10.006>
- Dubhashi, D., & Lappin, S. (2017). AI Dangers: Imagined and Real. *Communications of the ACM*, 60(2), 43–45. <http://doi.org/10.1145/2953876>
- Duncan, J. (1984). Selective Attention and the Organization of Visual Information. *Journal of Experimental Psychology: General*, 113(4), 501–17.
- Duncan, M. (2017). The Self Shows Up in Experience. *Review of Philosophy and Psychology*. <http://doi.org/10.1007/s13164-017-0355-2>
- Egeth, H. E., Leonard, C. J., & Leber, A. B. (2010). Why salience is not enough: reflections on top-down selection in vision. *Acta Psychologica*, 135(2), 130–139.

- <http://doi.org/10.1016/j.actpsy.2010.05.012>
- Eimer, M. (1996). The N2pc component as an indicator of attentional selectivity. *Electroencephalography and Clinical Neurophysiology*, 99(3), 225–234. [http://doi.org/10.1016/0013-4694\(96\)95711-9](http://doi.org/10.1016/0013-4694(96)95711-9)
- Eimer, M. (2015). The control of attention in visual search: Cognitive and neural mechanisms. *The Quarterly Journal of Experimental Psychology*, 68(12), 2437–2463.
- Engle, R. W., & Kane, M. J. (2004). Executive attention, working memory capacity, and a two-factor theory of cognitive control. *Psychology of Learning and Motivation*, 44, 145–199.
- Engle, R. W., Kane, M. J., & Tuholski, S. W. (1999). Individual Differences in Working Memory Capacity and What They Tell Us About Controlled Attention, General Fluid Intelligence, and Functions of the Prefrontal Cortex. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 102–134). Cambridge University Press.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16(1), 143–149.
- Eriksen, C. W., & Hoffman, J. E. (1972). Temporal and spatial characteristics of selective encoding from visual displays. *Perception & Psychophysics*, 12(2), 201–204. <http://doi.org/10.3758/BF03212870>
- Eriksen, C. W., & St James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, 40(4), 225–240. <http://doi.org/10.3758/BF03211502>
- Eriksson, J., Vogel, E. K., Lansner, A., Bergström, F., & Nyberg, L. (2015). Neurocognitive Architecture of Working Memory. *Neuron*, 88(1), 33–46. <http://doi.org/10.1016/j.neuron.2015.09.020>
- Falkenstein, L., & Grandi, G. (2017). Étienne Bonnot de Condillac. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter). Metaphysics Research Lab, Stanford University.
- Farrell, J. (2018). Higher-order theories of consciousness and what-it-is-like-ness. *Philosophical Studies*, 175(11), 2743–2761. <http://doi.org/10.1007/s11098-017-0980-8>
- Fazekas, P., & Nanay, B. (2018). Attention Is Amplification, Not Selection. *The British Journal for the Philosophy of Science*, axy065. Retrieved from <http://dx.doi.org/10.1093/bjps/axy065>
- Fazekas, P., & Nemeth, G. (2018). Dream experiences and the neural correlates of perceptual consciousness and cognitive access. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170356. <http://doi.org/10.1098/rstb.2017.0356>
- Fazekas, P., Nemeth, G., & Overgaard, M. (2019). White dreams are made of colours: What studying contentless dreams can teach about the neural basis of dreaming and conscious experiences. *Sleep Medicine Reviews*, 43, 84–91. <http://doi.org/https://doi.org/10.1016/j.smrv.2018.10.005>
- Fazekas, P., & Overgaard, M. (2018). Perceptual consciousness and cognitive access: an introduction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 373(1755), 20170340. <http://doi.org/10.1098/rstb.2017.0340>
- Fellrath, J., & Ptak, R. (2015). The role of visual saliency for the allocation of attention: Evidence from spatial neglect and hemianopia. *Neuropsychologia*, 73, 70–81. <http://doi.org/10.1016/j.neuropsychologia.2015.05.003>

- Fields, R. D. (2009). *The Other Brain: from dementia to schizophrenia, how new discoveries about the brain are revolutionizing medicine and science*. Simon & Schuster.
- Fields, R. D. (2013). Neuroscience: Map the other brain. *Nature*, 501(7465), 25–27. <http://doi.org/10.1038/501025a>
- Fine, C. (2010). From Scanner to Sound Bite: Issues in Interpreting and Reporting Sex Differences in the Brain. *Current Directions in Psychological Science*, 19(5), 280–283. <http://doi.org/10.1177/0963721410383248>
- Fisher, C. (n.d.). Prosochē: Illuminating the Path of the Prokopton. Retrieved November 29, 2017, from http://www.collegeofstoicphilosophers.org/show_book/PDF/ChrisFisher-Prosoche-IlluminatingthePathoftheProkopton
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews. Neuroscience*, 10(1), 48–58. <http://doi.org/10.1038/nrn2536>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... Vayena, E. (2018). AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689–707. <http://doi.org/10.1007/s11023-018-9482-5>
- Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press.
- Ford, M. (2013). Could Artificial Intelligence Create an Unemployment Crisis? *Communications of the ACM*, 56(7), 37–39. <http://doi.org/10.1145/2483852.2483865>
- Forman, R. K. (1999). What Does Mysticism Have to Teach Us About Consciousness? In *Models of the Self* (pp. 614–645). Imprint Academic (ProQuest Ebook Central).
- Förstl, H., & Kurz, A. (1999). Clinical features of Alzheimer's disease. *European Archives of Psychiatry and Clinical Neuroscience*, 249(6), 288–290.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27), 9673–8.
- Franconeri, S. L., Alvarez, G. A., & Enns, J. T. (2007). How Many Locations Can Be Selected at Once? *Journal of Experimental Psychology: Human Perception and Performance*, 33(5), 1003–1012. <http://doi.org/10.1037/0096-1523.33.5.1003>
- Frankfurt, H. (2003). Freedom of the Will and the Concept of a Person. In G. Watson (Ed.), *Free Will* (2nd ed., pp. 322–336). Oxford University Press.
- Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. (2014). Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception. *The Journal of Neuroscience*, 34(5), 1738–1747. Retrieved from <http://www.jneurosci.org/content/34/5/1738.abstract>
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301. <http://doi.org/10.1016/j.tics.2009.04.005>
- Fulkerson, M. (2014). Rethinking the senses and their interactions: the case for sensory pluralism. *Frontiers in Psychology*, 5(1426), 1–14. <http://doi.org/10.3389/fpsyg.2014.01426>
- Galin, D. (1994). The Structure of Awareness: Contemporary Applications of William James' Forgotten Concept of "The Fringe". *The Journal of Mind and Behavior*, 15(4), 375–401.
- Gallagher, S. (1998). *The Inordinance of Time*. Northwestern University Press.

- Gallagher, S. (2004). Agency, ownership and alien control in schizophrenia. In P. Bovet, J. Parnas, & D. Zahavi (Eds.), *Interdisciplinary Perspectives on Self-Consciousness* (pp. 89–104). John Benjamins Publishers.
- Gallagher, S. (Ed.). (2011). *The Oxford Handbook of the Self*. Oxford University Press.
- Gallagher, S., & Shear, J. (Eds.). (1999). *Models of the Self*. Imprint Academic.
- Galvin, S. J., O'Shea, R. P., Squire, A. M., & Govan, D. G. (1997). Sharpness overconstancy in peripheral vision. *Vision Research*, 37(15), 2035–2039.
[http://doi.org/https://doi.org/10.1016/S0042-6989\(97\)00016-3](http://doi.org/https://doi.org/10.1016/S0042-6989(97)00016-3)
- Ganeri, J. (2017). *Attention, not self*. Oxford University Press.
- Gangemi, A., Bourgeois-Gironde, S., & Mancini, F. (2014). Feelings of error in reasoning – in search of a phenomenon. *Thinking & Reasoning*, 21(4), 383–396.
<http://doi.org/10.1080/13546783.2014.980755>
- Gayet, S., Paffen, C. L. E., & der Stigchel, S. Van. (2018). Visual Working Memory Storage Recruits Sensory Processing Areas. *Trends in Cognitive Sciences*, 22(3), 189–190.
<http://doi.org/10.1016/j.tics.2017.09.011>
- Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: bridging selective attention and working memory. *Trends in Cognitive Sciences*, 16(2), 129–135.
<http://doi.org/10.1016/j.tics.2011.11.014>
- Gelbard-Sagiv, H., Faivre, N., Mudrik, L., & Koch, C. (2016). Low-level awareness accompanies “unconscious” high-level processing during continuous flash suppression. *Journal of Vision*, 16(1), 1–16. <http://doi.org/10.1167/16.1.3>
- Giacino, J., Ashwal, S., Childs, N., Cranford, R., Jennett, B., Katz, D., ... Zasler, N. (2002). The minimally conscious state: Definition and diagnostic criteria. *Neurology*, 58(3), 349–353. <http://doi.org/10.1212/WNL.58.3.506>
- Godfrey-Smith, P. (2016). Animal Evolution and the Origins of Experience. In D. Livingstone Smith (Ed.), *How Biology Shapes Philosophy: New Foundations for Naturalism* (forthcomin, pp. 51–71). Cambridge University Press.
- Godman, D. (Ed.). (1985). *Be As You Are: The Teachings of Sri Ramana Maharshi*. Penguin Arkana.
- González-Álvarez, J., & Cervera-Crespo, T. (2017). Research production in high-impact journals of contemporary neuroscience: A gender analysis. *Journal of Informetrics*, 11(1), 232–243. <http://doi.org/10.1016/j.joi.2016.12.007>
- Grammer, K., Fink, B., Møller, A. P., & Thornhill, R. (2003). Darwinian aesthetics: sexual selection and the biology of beauty. *Biological Reviews*, 78(3), 385–407.
<http://doi.org/10.1017/S1464793102006085>
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27.
- Greenwood, J. A., Bex, P. J., & Dakin, S. C. (2010). Crowding Changes Appearance. *Current Biology*, 20(6), 496–501. <http://doi.org/https://doi.org/10.1016/j.cub.2010.01.023>
- Groes, S. (Ed.). (2016). *Memory in the Twenty-First Century: New Critical Perspectives from the Arts, Humanities, and Sciences*. London: Palgrave Macmillan UK.
<http://doi.org/10.1057/9781137520586>
- Gross, S. (2018). Perceptual consciousness and cognitive access from the perspective of capacity-unlimited working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170343. <http://doi.org/10.1098/rstb.2017.0343>

- Gross, S., & Flombaum, J. (2017). Does Perceptual Consciousness Overflow Cognitive Access? The Challenge from Probabilistic, Hierarchical Processes. *Mind & Language*, 32(3), 358–391.
- Guidolin, D., Albertin, G., Guescini, M., Fuxe, K., & Agnati, L. F. (2011). Central nervous system and computation. *The Quarterly Review of Biology*, 86(4), 265–85. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/22384748>
- Haladjian, H. H., & Montemayor, C. (2015). On the evolution of conscious attention. *Psychonomic Bulletin & Review*, 22(3), 595–613. <http://doi.org/DOI 10.3758/s13423-014-0718-y>
- Hall, J. E. (2015). *Guyton and Hall Textbook of Medical Physiology*. Elsevier.
- Hanczakowski, M., Zawadzka, K., Collie, H., & Macken, B. (2017). Metamemory in a Familiar Place: The Effects of Environmental Context on Feeling of Knowing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(1), 59–71. <http://doi.org/10.1037/xlm0000292>
- Hanoğlu, L., Özkara, Ç., Yalçiner, B., Nani, A., & Cavanna, A. E. (2014). Epileptic qualia and self-awareness: A third dimension for consciousness. *Epilepsy & Behavior*, 30, 62–65. <http://doi.org/https://doi.org/10.1016/j.yebeh.2013.09.010>
- Hardcastle, V. G. (1997). Attention versus consciousness: a distinction with a difference. *Bulletin of the Japanese Cognitive Science Society*, 4(3), 56–66.
- Hardman, K., & Cowan, N. (2015). Remembering Complex Objects in Visual Working Memory: Do Capacity Limits Restrict Objects or Features? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(2), 325–347. <http://doi.org/10.1037/xlm0000031>
- Harrison, A. H., & Connolly, J. F. (2013). Finding a way in: A review and practical evaluation of fMRI and EEG for detection and assessment in disorders of consciousness. *Neuroscience and Biobehavioral Reviews*, 37(8), 1403–1419. <http://doi.org/10.1016/j.neubiorev.2013.05.004>
- Hassin, R. R. (2013). Yes It Can: On the Functional Abilities of the Human Unconscious. *Perspectives on Psychological Science*, 8(2), 195–207. <http://doi.org/10.1177/1745691612460684>
- Hassin, R. R., Bargh, J. A., Engell, A. D., & McCulloch, K. C. (2009). Implicit working memory. *Consciousness and Cognition*, 18(3), 665–678. <http://doi.org/10.1016/j.concog.2009.04.003>
- Haun, A. M., Tononi, G., Koch, C., & Tsuchiya, N. (2017). Are we underestimating the richness of visual experience? *Neuroscience of Consciousness*, 2017(1), Niw023. <http://doi.org/10.1093/nc/niw023>
- Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, 7(11), 498–504. <http://doi.org/10.1016/j.tics.2003.09.006>
- Herzog, M. H., Kammer, T., & Scharnowski, F. (2016). Time Slices: What Is the Duration of a Percept? *PLoS Biology*, 14(4). <http://doi.org/10.1371/journal.pbio.1002433>
- Hilgard, E. R., Hilgard, J. R., Macdonald, H., Morgan, A. H., & Johnson, L. S. (1978). Covert pain in hypnotic analgesia: Its reality as tested by the real-simulator design. *Journal of Abnormal Psychology*, 87(6), 655–663.
- Hine, R. (2010). Attention as Experience: Through “Thick” & “Thin.” *Journal of Consciousness Studies*, 17(9–10), 202–220.
- Hine, R. (2015). Paying Attention to Consciousness. *Journal of Consciousness Studies*, 22(5–

6), 52–69.

- Hine, R. J. (2010). *Paying Attention to Consciousness*. University of Connecticut.
- Hohwy, J. (2012). Attention and Conscious Perception in the Hypothesis Testing Brain. *Frontiers in Psychology*, 3, 96. <http://doi.org/10.3389/fpsyg.2012.00096>
- Holcombe, A. O., & Cavanagh, P. (2001). Early binding of feature pairs for visual perception. *Nature Neuroscience*, 4, 127.
- Holcombe, A. O., Linares, D., & Vaziri-Pashkam, M. (2011). Perceiving Spatial Relations via Attentional Tracking and Shifting. *Current Biology*. <http://doi.org/10.1016/j.cub.2011.05.031>
- Holmes, A. J., & Patrick, L. M. (2018). The Myth of Optimality in Clinical Neuroscience. *Trends in Cognitive Sciences*, 22(3), 241–257. <http://doi.org/10.1016/j.tics.2017.12.006>
- Holtzman, G. S. (2016). Rejecting Beliefs, or Rejecting Believers? On the Importance and Exclusion of Women in Philosophy. *Hypatia*, 31(2), 293–312.
- Horgan, T. (2012). Introspection about Phenomenal Consciousness: Running the Gamut from Infallibility to Impotence. In D. Smithies & D. Stoljar (Eds.), *Introspection and Consciousness*. Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199744794.003.0015>
- Hsieh, P.-J., & Colas, J. T. (2012). Awareness Is Necessary for Extracting Patterns in Working Memory but Not for Directing Spatial Attention. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5), 1085–1090. <http://doi.org/10.1037/a0028345>
- Hung, S.-M., & Hsieh, P.-J. (2015). Syntactic Processing in the Absence of Awareness and Semantics. *Journal of Experimental Psychology: Human Perception and Performance*, 41(5), 1376–1384. <http://doi.org/10.1037/xhp0000094>
- Hutchison, K., & Jenkins, F. (2013). *Women in philosophy: What needs to change?* Oxford University Press.
- Hyslop, A. (2018). Other Minds. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 201). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/win2018/entries/other-minds/>
- Intriligator, J., & Cavanagh, P. (2001). The Spatial Resolution of Visual Attention. *Cognitive Psychology*, 43(3), 171–216. <http://doi.org/https://doi.org/10.1006/cogp.2001.0755>
- Irving, Z. C. (2016). Mind-wandering is unguided attention: accounting for the “purposeful” wanderer. *Philosophical Studies*, 173(2), 547–571. <http://doi.org/10.1007/s11098-015-0506-1>
- Irving, Z. C., & Thompson, E. (2017). The philosophy of mind-wandering. Retrieved from https://www.academia.edu/32057733/The_Philosophy_of_Mind-Wandering
- Iwasaki, S. (1993). Spatial attention and two modes of visual consciousness. *Cognition*, 49(3), 211–233.
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. MIT Press.
- Jacob, J., Jacobs, C., & Silvanto, J. (2015). Attention, working memory, and phenomenal experience of WM content: Memory levels determined by different types of top-down modulation. *Frontiers in Psychology*, 6(1603), 1–7. <http://doi.org/10.3389/fpsyg.2015.01603>
- Jacobs, C., & Silvanto, J. (2015). How is working memory content consciously experienced? The ‘conscious copy’ model of WM introspection. *Neuroscience & Biobehavioral*

Reviews, 55, 510–519.

- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087), 1454–1458.
- James, W. (1890). *The Principles of Psychology*. Henry Holt and Co.
- Jarrold, C., Tam, H., Baddeley, A. D., & Harvey, C. (2011). How Does Processing Affect Storage in Working Memory Tasks? Evidence for Both Domain-General and Domain-Specific Effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(3), 688–705. <http://doi.org/10.1037/a0022527>
- Jaynes, J. (1976). *The origin of consciousness in the breakdown of the bicameral mind*. Boston: Houghton Mifflin.
- Jennings, C. D. (2012). The subject of attention. *Synthese*, 189, 535–554. <http://doi.org/10.1007/s11229-012-0164>
- Jennings, C. D. (2015). Consciousness Without Attention. *Journal of the American Philosophical Association*, 1(02), 276–295. <http://doi.org/10.1017/apa.2014.14>
- Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*, 103(45), 17048–17052. <http://doi.org/10.1073/pnas.0605678103>
- Johanson, M., Revonsuo, A., Chaplin, J., & Wedlund, J.-E. (2003). Level and contents of consciousness in connection with partial epileptic seizures. *Epilepsy & Behavior*, 4(3), 279–285. [http://doi.org/10.1016/S1525-5050\(03\)00106-9](http://doi.org/10.1016/S1525-5050(03)00106-9)
- Jonides, J., Lewis, R. L., Nee, D. E., Lustig, C. A., Berman, M. G., & Moore, K. S. (2008). The Mind and Brain of Short-Term Memory. *Annual Review of Psychology*, 59, 193–224.
- Joyston-Bechal, M. (1966). The Clinical Features and Outcome of Stupor. *British Journal of Psychiatry*, 112(491), 967–981. <http://doi.org/10.1192/bjp.112.491.967>
- Kaldas, A. (2015). *The One and the Many of Experience: The tension between mereology and holism in Tim Bayne's account of the unity of consciousness*. Macquarie University. Retrieved from <http://minerva.mq.edu.au:8080/vital/access/manager/Repository/mq:44784>
- Kanai, R., Tsuchiya, N., & Verstraten, F. A. (2006). The Scope and Limits of Top-Down Attention in Unconscious Visual Processing. *Current Biology*, 16(23), 2332–2336. <http://doi.org/10.1016/j.cub.2006.10.001>
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review*, 9(4), 637–671. <http://doi.org/10.3758/BF03196323>
- Kanwisher, N., & Wojciulik, E. (2000). Visual attention: insights from brain imaging. *Nature Reviews Neuroscience*, 1(2), 91–100. <http://doi.org/10.1038/35039043>
- Katusic, M. Z., Voigt, R. G., Colligan, R. C., Weaver, A. L., Homan, K. J., & Barbaresi, W. J. (2011). Attention-Deficit Hyperactivity Disorder in Children With High Intelligence Quotient: Results From a Population-Based Study. *Journal of Developmental & Behavioral Pediatrics*, 32(2), 103–109. <http://doi.org/10.1097/DBP.0b013e318206d700>
- Kaukua, J. (2015). *Self-Awareness in Islamic Philosophy: Avicenna and Beyond*. Cambridge University Press.
- Keller, A. (2011). Attention and olfactory consciousness. *Frontiers in Psychology*, 2(380), 1–13. <http://doi.org/10.3389/fpsyg.2011.00380>

- Kennett, J., & Fine, C. (2009). Will the Real Moral Judgment Please Stand Up? *Ethical Theory and Moral Practice*, 12(1), 77–96. <http://doi.org/10.1007/s10677-008-9136-4>
- Kennett, J., & Matthews, S. (2014). The Unity and Disunity of Agency. *Philosophy, Psychiatry, & Psychology*, 10(4), 305–312. <http://doi.org/10.1353/ppp.2004.0020>
- Kentridge, R. W. (2011). Attention Without Awareness: A Brief Review. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 228–246). Oxford University Press.
- Kentridge, R. W., & Heywood, C. A. (2001). Attention and alerting: Cognitive processes spared in blindsight. In B. De Gelder, E. De Haan, & C. Heywood (Eds.), *Varieties of Unconscious Processing: New Findings and Models*. Oxford University Press.
- Kentridge, R. W., Heywood, C., & Weiskrantz, L. (1999). Attention without awareness in blindsight. *Proceedings of the Royal Society B: Biological Sciences*, 266, 1805–1811. <http://doi.org/10.1098/rspb.1999.0850>
- Kentridge, R. W., Heywood, C., & Weiskrantz, L. (2004). Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia*, 42(6), 831–835. <http://doi.org/10.1016/j.neuropsychologia.2003.11.001>
- Kentridge, R. W., Nijboer, T., & Heywood, C. (2008). Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia*, 46(3), 864–869. <http://doi.org/10.1016/j.neuropsychologia.2007.11.036>
- Key, B., Arlinghaus, R., & Browman, H. I. (2016). Insects cannot tell us anything about subjective experience or the origin of consciousness. *Proceedings of the National Academy of Sciences*, 113(27), E3813. <http://doi.org/10.1073/pnas.1606835113>
- Kihlstrom, J. F. (1997). The Psychological Unconscious. In C. Prendergast (Ed.), *Napoleon and History Painting: Antoine-Jean Gros's La Bataille D'Eylau* (pp. 424–442). Oxford University Press.
- Kihlstrom, J. F. (2007). Consciousness in Hypnosis. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 445–479). Cambridge University Press.
- Kimchi, R., Devyatko, D., & Sabary, S. (2018). Can perceptual grouping unfold in the absence of awareness? Comparing grouping during continuous flash suppression and sandwich masking. *Consciousness and Cognition*, 60, 37–51. <http://doi.org/10.1016/j.concog.2018.02.009>
- Klein, C. (2017). Consciousness, intention, and command-following in the vegetative state. *British Journal for the Philosophy of Science*, 68(1), 27–54. <http://doi.org/10.1093/bjps/axv012>
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, 17(5), 307–321.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22. <http://doi.org/10.1016/j.tics.2006.10.012>
- Koch, C., & Tsuchiya, N. (2008). Response to Mole: Subjects can attend to completely invisible objects. *Trends in Cognitive Sciences*, 12(2), 44–45. <http://doi.org/10.1016/j.tics.2007.11.002>
- Koch, C., & Tsuchiya, N. (2012). Attention and consciousness: related yet different. *Trends in Cognitive Sciences*, 16(2), 103–105. <http://doi.org/10.1016/j.tics.2011.11.012>
- Koivisto, M., Kainulainen, P., & Revonsuo, A. (2009). The relationship between awareness and attention: Evidence from ERP responses. *Neuropsychologia*, 47(13), 2891–2899.

- <http://doi.org/10.1016/j.neuropsychologia.2009.06.016>
- Koivisto, M., & Revonsuo, A. (2007). Electrophysiological correlates of visual consciousness and selective attention. *Neuroreport*, 18(8), 753–6.
<http://doi.org/10.1097/WNR.0b013e3280c143c8>
- Koivisto, M., Revonsuo, A., & Salminen, N. (2005). Independence of visual awareness from attention at early processing stages. *Neuroreport*, 16(8), 817–821.
<http://doi.org/10.1097/00001756-200505310-00008>
- Koizumi, A., Maniscalco, B., & Lau, H. (2015). Does perceptual confidence facilitate cognitive control? *Attention, Perception, and Psychophysics*, 77(4), 1295–1306.
<http://doi.org/10.3758/s13414-015-0843-3>
- Kondziella, D. (2017). Roald Dahl and the complete locked-in syndrome: “Cold dead body, living brain.” *Journal of the Neurological Sciences*, 379, 276–278.
<http://doi.org/https://doi.org/10.1016/j.jns.2017.06.033>
- Kondziella, D., Friberg, C. K., Frokjaer, V. G., Fabricius, M., & Møller, K. (2016). Preserved consciousness in vegetative and minimal conscious states: systematic review and meta-analysis. *Journal of Neurology, Neurosurgery, & Psychiatry*, 87(5), 485–492.
- Koralus, P. (2014). The Erotetic Theory of Attention: Questions, Focus and Distraction. *Mind & Language*, 29(1), 26–50.
- Koriat, A., Goldsmith, M., & Pansky, A. (2000). Toward a Psychology of Memory Accuracy. *Annual Review of Psychology*, 51, 481–537.
- 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>
- Kozuch, B., & Kriegel, U. (2015). Correlation, causation, constitution: On the interplay between the science and philosophy of consciousness. In S. M. Miller (Ed.), *The Constitution of Phenomenal Consciousness* (pp. 400–417). John Benjamins.
- Kriegel, U. (2004). The functional role of consciousness: A phenomenological approach. *Phenomenology and the Cognitive Sciences*, 3(2), 171–193.
<http://doi.org/10.1023/B:PHEN.0000040833.23356.6a>
- Kriegel, U. (2007). The phenomenologically manifest. *Phenomenology and the Cognitive Sciences*, 6(1–2), 115–136. <http://doi.org/10.1007/s11097-006-9029-8>
- Kvanvig, J. (2007). Propositionalism and the Metaphysics of Experience. *Philosophical Issues*, 17, 165–178.
- Lähteenmäki, V. (2010). Cudworth on Types of Consciousness. *British Journal for the History of Philosophy*, 18(1), 9–34. <http://doi.org/10.1080/09608780903422206>
- Lamme, V. A. (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. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks*, 17(5–6), 861–872.
<http://doi.org/10.1016/j.neunet.2004.02.005>
- Lamme, V. A. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501. <http://doi.org/10.1016/j.tics.2006.09.001>

- Lamme, V. A. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, 1(3), 204–220. <http://doi.org/10.1080/17588921003731586>
- Lamme, V. A. (2018). Challenges for theories of consciousness: seeing or knowing, the missing ingredient and how to deal with panpsychism. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170344. <http://doi.org/10.1098/rstb.2017.0344>
- Landman, R., Spekreijse, H., & Lamme, V. A. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2), 149–164. [http://doi.org/https://doi.org/10.1016/S0042-6989\(02\)00402-9](http://doi.org/https://doi.org/10.1016/S0042-6989(02)00402-9)
- Lavie, N. (2010). Attention, distraction, and cognitive control under load. *Current Directions in Psychological Science*, 19(3), 143–148.
- Le Poidevin, R. (2015). The Experience and Perception of Time. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Summer, Vol. Summer). Stanford University. Retrieved from <http://plato.stanford.edu/archives/sum2015/entries/time-experience/>
- Lee, G. (2014). Experiences and their Parts. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 287–321). MIT Press.
- Letheby, C., & Gerrans, P. (2017). Self unbound: ego dissolution in psychedelic experience. *Neuroscience of Consciousness*, 2017(1), 1–11. <http://doi.org/10.1093/nc/nix016>
- Lev, M., & Barkai, R. (2016). Elephants are people, people are elephants: Human–proboscideans similarities as a case for cross cultural animal humanization in recent and Paleolithic times. *Quaternary International*, 406, 239–245. <http://doi.org/10.1016/j.quaint.2015.07.005>
- Lewandowsky, S., & Oberauer, K. (2015). Rehearsal in Serial Recall: An Unworkable Solution to the Nonexistent Problem of Decay. *Psychological Review*, 122(4), 674–699. <http://doi.org/10.1037/a0039684>
- Lewis, P. A. (2013). *The Secret World of Sleep: The Surprising Science of the Mind at Rest*. Palgrave Macmillan.
- Li, D., Christ, S. E., & Cowan, N. (2014). Domain-general and domain-specific functional networks in working memory. *NeuroImage*, 102, 646–656. <http://doi.org/10.1016/j.neuroimage.2014.08.028>
- Li, F. F., Vanrullen, R., Koch, C., & Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences of the United States of America*, 99(14), 9596–9601.
- Ling, S. (2012). Attention alters appearance. *Journal of Vision*, 12(9), 1387.
- Logie, R. H., & Cowan, N. (2015). Perspectives on working memory: introduction to the special issue. *Memory & Cognition*, 43(3), 315–324. <http://doi.org/10.3758/s13421-015-0510-x>
- Long, A. (1986). *Hellenistic Philosophy: Stoics, Epicureans, Sceptics* (2nd ed.). University of California Press.
- Long, G. M. (1980). Iconic memory: A review and critique of the study of short-term visual storage. *Psychological Bulletin*, 88(3), 785–820.
- Louth, A. (2007). *The Origins of the Christian Mystical Tradition: From Plato to Denys* (2nd ed.). Oxford University Press.
- Lovett, M. C., Reder, L. M., & Lebiere, C. (1999). Modeling Working Memory in a Unified Architecture: An ACT-R Perspective. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 135–182).

- Cambridge University Press.
- Lowe, E. (2002). *A Survey of Metaphysics*. Oxford University Press.
- Luck, S. J., & Hillyard, S. A. (1994). Spatial Filtering During Visual Search: Evidence From Human Electrophysiology. *Journal of Experimental Psychology: Human Perception and Performance*, 20(5), 1000–1014.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Lukens, H. T. (1896). Preliminary report on the learning of language. *The Journal of Genetic Psychology*, 3(3), 424–460.
- Lundqvist, M., Herman, P., & Miller, E. K. (2018). Working Memory: Delay Activity, Yes! Persistent Activity? Maybe Not. *The Journal of Neuroscience*, 38(32), 7013–7019. <http://doi.org/10.1523/JNEUROSCI.2485-17.2018>
- Lutz, A., Dunne, J. D., & Davidson, R. J. (2007). Meditation and the Neuroscience of Consciousness: An Introduction. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness* (pp. 499–551). Cambridge University Press.
- Lutz, A., Slagter, H. A., Dunne, J. D., & Davidson, R. J. (2008). Attention regulation and monitoring in meditation. *Trends in Cognitive Sciences*, 12(4), 163–169. <http://doi.org/https://doi.org/10.1016/j.tics.2008.01.005>
- Lycan, W. G. (1996). *Consciousness and experience*. Cambridge, Mass.: MIT Press.
- Lyons, J. (2017). Epistemological Problems of Perception. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2017/entries/perception-episprob/>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17(3), 347–356. <http://doi.org/10.0.4.14/nn.3655>
- Mack, A., & Clarke, J. (2012). Gist perception requires attention. *Visual Cognition*, 20(3), 300–327. <http://doi.org/10.1167/10.7.191>
- Mack, A., Clarke, J., & Erol, M. (2018). Attention, expectation and iconic memory: A reply to Aru and Bachmann (2017). *Consciousness and Cognition*, 59, 60–63. <http://doi.org/10.1016/j.concog.2017.10.001>
- Mack, A., & Rock, I. (1998). *Inattentional Blindness*. MIT Press.
- Maehara, Y., & Saito, S. (2007). The Relationship between Processing and Storage in Working Memory Span: Not Two Sides of the Same Coin. *Journal of Memory and Language*, 56(2), 212–228. <http://doi.org/10.1016/j.jml.2006.07.009>
- Magnussen, S. (2009). Implicit visual working memory. *Cognition and Neurosciences*, 50, 535–542. <http://doi.org/10.1111/j.1467-9450.2009.00783.x>
- Malyshev, E. (2013). Two Metaphors of Memory in Early Modern Philosophy. *Problemas*, 84, 36–45.
- Mangan, B. (2001). Sensation's Ghost: The Non-Sensory "Fringe" of Consciousness. *Psyche*, 7(18). Retrieved from <http://psyche.cs.monash.edu.au/v7/psyche-7-18-mangan.html>
- Mangan, B. (2007). Cognition, Fringe Consciousness, and the Legacy of William James. In M. Velmans & S. Schneider (Eds.), *The Blackwell Companion to Consciousness* (pp. 673–685). Blackwell. <http://doi.org/doi:10.1002/9780470751466.ch53>
- Maniscalco, B., & Lau, H. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task.

- Neuroscience of Consciousness*, 1–13. <http://doi.org/10.1093/nc/niv002>
- Maniscalco, B., & Lau, H. (2016). The signal processing architecture underlying subjective reports of sensory awareness. *Neuroscience of Consciousness*, 1–17. <http://doi.org/10.1093/nc/niw002>
- Marr, D. (1982). *Vision*. Freeman.
- Masrour, F. (2011). Is perceptual phenomenology thin? *Philosophy and Phenomenological Research*, 83(2), 366–397.
- Matthews, J., Schröder, P., Kaunitz, L., van Boxtel, J. J. A., & Tsuchiya, N. (2018). Conscious access in the near absence of attention: critical extensions on the dual-task paradigm. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170352. <http://doi.org/10.1098/rstb.2017.0352>
- May, J. (2004). An information processing view of fringe consciousness. *Psyche*, 10(1), 1–9.
- McCarroll, C. J. (2018). Looking at the self: perspectival memory and personal identity. *Philosophical Explorations*, 1–21. <http://doi.org/10.1080/13869795.2018.1562087>
- McCarthy, J., & Hayes, P. (1969). Some philosophical problems from the stand-point of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4* (pp. 463–502). Edinburgh University Press.
- Menary, R. (2007). Writing as thinking. *Language Sciences*, 29(5), 621–632. <http://doi.org/10.1016/j.langsci.2007.01.005>
- Menary, R. (2009). Intentionality and Consciousness. In W. Banks (Ed.), *Encyclopedia of Consciousness* (1st ed., Vol. 1, pp. 417–430). Academic Press.
- Menary, R. (2010). Introduction to the special issue on 4E cognition. *Phenomenology and the Cognitive Sciences*, 9(4), 459–463. <http://doi.org/10.1007/s11097-010-9187-6>
- Menary, R. (2018). Keeping Track with Things. In J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos, & D. Pritchard (Eds.), *Extended Epistemology* (pp. 305–330). Oxford University Press. <http://doi.org/10.1093/oso/9780198769811.003.0016>
- Merikle, P. M. (1992). Perception Without Awareness: Critical Issues. *American Psychologist*, 47(6), 792–795.
- Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, 30(1), 63–81. <http://doi.org/10.1017/S0140525X07000891>
- Metzinger, T. (2003). *Being no one: the self-model theory of subjectivity*. Cambridge, Mass.: Cambridge, Mass. : MIT Press.
- Metzinger, T. (2005). Precis: Being no one. *Psyche*, 11, 1–35.
- Michaelian, K., & Sutton, J. (2017). Memory. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Summer). Retrieved from <https://plato.stanford.edu/archives/sum2017/entries/memory/>
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97. <http://doi.org/10.1037/h0043158>
- Miller, G. A. (1962). *Psychology, the science of mental life*. Hutchinson.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. Henry Holt and Co.
- Mischel, W. (2008). The toothbrush problem. *APS Observer*, 21(11). Retrieved from

- <https://www.psychologicalscience.org/observer/the-toothbrush-problem>
- Mitchell, S. D. (2000). Dimensions of Scientific Law. *Philosophy of Science*, 67(2).
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The Unity and Diversity of Executive Functions and Their Contributions to Complex Frontal Lobe Tasks: A Latent Variable Analysis. *Cognitive Psychology*, 41(1), 49–100. <http://doi.org/10.1006/cogp.1999.0734>
- Miyake, A., & Shah, P. (Eds.). (1999). *Models of Working Memory: mechanisms of active maintenance and executive control*. Cambridge University Press.
- Mogensen, J., & Overgaard, M. (2018). Reorganization of the connectivity between elementary functions as a common mechanism of phenomenal consciousness and working memory: from functions to strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170346. <http://doi.org/10.1098/rstb.2017.0346>
- Mole, C. (2008a). Attention and consciousness. *Journal of Consciousness Studies*, 15(4), 86–104.
- Mole, C. (2008b). Attention in the absence of consciousness? *Trends in Cognitive Sciences*, 12(2), 44. <http://doi.org/doi:10.1016/j.tics.2007.11.001>
- Mole, C. (2011a). *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780195384529.001.0001>
- Mole, C. (2011b). The Metaphysics of Attention. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 60–77). Oxford University Press.
- Mole, C. (2013). Attention. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Fall, 2013). Stanford University. Retrieved from <http://plato.stanford.edu/archives/fall2013/entries/attention/>
- Mole, C. (2014). Attention to Unseen Objects. *Journal of Consciousness Studies*, 21(11–12), 41–56.
- Mole, C. (2016). *The Unexplained Intellect: Complexity, Time, and the Metaphysics of Embodied Thought*. Taylor and Francis.
- Mole, C., Smithies, D., & Wu, W. (Eds.). (2011a). *Attention: Philosophical and Psychological Essays*. Oxford University Press.
- Mole, C., Smithies, D., & Wu, W. (2011b). Introduction. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. xi–xx). Oxford University Press.
- Monaco, F., Mula, M., & Cavanna, A. E. (2005). Consciousness, epilepsy, and emotional qualia. *Epilepsy & Behavior*, 7(2), 150–160. <http://doi.org/https://doi.org/10.1016/j.yebeh.2005.05.018>
- Montaser-Kouhsari, L., & Rajimehr, R. (2004). Attentional modulation of adaptation to illusory lines. *Journal of Vision*, 4(6), 434–444. <http://doi.org/10.1167/4.6.3>
- Montemayor, C., & Haladjian, H. H. (2015). *Consciousness, Attention, and Conscious Attention*. MIT Press.
- Moore, T., Armstrong, K., & Fallah, M. (2003). Visuomotor Origins of Covert Spatial Attention. *Neuron*, 40(4), 671–683. [http://doi.org/10.1016/S0896-6273\(03\)00716-5](http://doi.org/10.1016/S0896-6273(03)00716-5)
- Moors, P., & Hesselmann, G. (2018). A critical reexamination of doing arithmetic nonconsciously. *Psychonomic Bulletin & Review*, 25(1), 472–481.
- Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of

- instructions. *Quarterly Journal of Experimental Psychology*, 11, 56–60.
- Morelli, G. (n.d.). Mindfulness as Known by the Church Fathers. Retrieved November 29, 2017, from <http://ww1.antiochian.org/mindfulness-known-church-fathers>
- Morin, A. (2006). Levels of consciousness and self-awareness: A comparison and integration of various neurocognitive views. *Consciousness and Cognition*, 15(2), 358–371. <http://doi.org/https://doi.org/10.1016/j.concog.2005.09.006>
- Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences*, 18(9), 488–496. <http://doi.org/10.1016/j.tics.2014.04.009>
- Munsell, O. (1873). *Psychology: The Science of the Mind*. D. Appleton.
- Myers, N. E., Stokes, M. G., & Nobre, A. C. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, 21(6), 449–461. <http://doi.org/10.1016/j.tics.2017.03.010>
- Naccache, L. (2006). Is She Conscious? *Science*, 313(5792), 1395–1396. <http://doi.org/10.126/science.1132881>
- Naccache, L. (2018). Why and how access consciousness can account for phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170357. <http://doi.org/10.1098/rstb.2017.0357>
- Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious Masked Priming Depends on Temporal Attention. *Psychological Science*, 13(5), 416–424.
- Nagel, T. (1974). What Is It Like to Be a Bat? *Philosophical Review*, 83(4), 435–450. <http://doi.org/10.2307/2183914>
- Naghavi, H. R., & Nyberg, L. (2005). Common fronto-parietal activity in attention, memory, and consciousness: Shared demands on integration? *Consciousness and Cognition*, 14(2), 390–425. <http://doi.org/10.1016/j.concog.2004.10.003>
- Naselaris, T., Bassett, D. S., Fletcher, A. K., Kording, K., Kriegeskorte, N., Nienborg, H., ... Kay, K. (2018). Cognitive Computational Neuroscience: A New Conference for an Emerging Discipline. *Trends in Cognitive Sciences*, 22(5), 365–367. <http://doi.org/10.1016/j.tics.2018.02.008>
- Nash, J. D., & Newberg, A. (2013). Toward a unifying taxonomy and definition for meditation. *Frontiers in Psychology*, 4(806), 1–18. <http://doi.org/10.3389/fpsyg.2013.00806>
- Neander, K. (1998). The division of phenomenal labor: A problem for representational theories of consciousness. *Philosophical Perspectives*, 32(S12), 411–434.
- Neisser, U. (1967). *Cognitive Psychology*. Prentice-Hall.
- Nemrodov, D., Niemeier, M., Patel, A., & Nestor, A. (2018). The Neural Dynamics of Facial Identity Processing: Insights from EEG-Based Pattern Analysis and Image Reconstruction. *ENeuro*, 5(1), 1–17. <http://doi.org/10.1523/ENEURO.0358-17.2018>
- Newell, A., & Simon, H. (1956). *The logic theory machine: A complex information processing system*. Rand Corp.
- Nielsen, T., O'Reilly, C., Carr, M., Dumel, G., Godin, I., Solomonova, E., ... Paquette, T. (2015). Overnight improvements in two REM sleep-sensitive tasks are associated with both REM and NREM sleep changes, sleep spindle features, and awakenings for dream recall. *Neurobiology of Learning and Memory*, 122, 88–97. <http://doi.org/https://doi.org/10.1016/j.nlm.2014.09.007>
- Nikulin, D. (Ed.). (2015). *Memory: A History*. Oxford University Press.
- Norman, D., & Shallice, T. (1986). Attention to action: Willed and automatic control of

- behavior. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.), *Consciousness and Self-Regulation* (pp. 1–18). Plenum Press.
- Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2015). Exogenous attention to unseen objects? *Consciousness and Cognition*, 35, 319–329.
<http://doi.org/10.1016/j.concog.2015.02.015>
- Normore, C. (2016). “Causa sui”: Awareness and Choice in the Constitution of the Self. In J. Kaukua & T. Ekenberg (Eds.), *Subjectivity and Selfhood in Medieval and Early Modern Philosophy* (pp. 91–107). Cham: Springer International Publishing.
http://doi.org/10.1007/978-3-319-26914-6_7
- Norris, D. (2017). Short-Term Memory and Long-Term Memory are Still Different. *Psychological Bulletin*, 143(9), 992–1009.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–973.
- Oatley, K., & Djikic, M. (2008). Writing as Thinking. *Review of General Psychology*, 12(1), 9–27. <http://doi.org/10.1037/1089-2680.12.1.9>
- Oberauer, K. (2002). Access to Information in Working Memory: Exploring the Focus of Attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 411–421. <http://doi.org/10.1037/0278-7393.28.3.411>
- Oberauer, K. (2009). Design for a working memory. *Psychology of Learning and Motivation*, 51, 45–100.
- Oberauer, K., Farrell, S., Jarrold, C., & Lewandowsky, S. (2016). What Limits Working Memory Capacity? *Psychological Bulletin*, 142(7), 758–799.
<http://doi.org/10.1037/bul0000046>
- Oberauer, K., & Lewandowsky, S. (2011). Modeling working memory: a computational implementation of the Time-Based Resource-Sharing theory. *Psychonomic Bulletin & Review*, 18(1), 10–45. <http://doi.org/10.3758/s13423-010-0020-6>
- Oberauer, K., Süß, H.-M., Wilhelm, O., & Wittman, W. W. (2003). The multiple faces of working memory: Storage, processing, supervision, and coordination. *Intelligence*, 31(2), 167–193.
- Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170345.
<http://doi.org/10.1098/rstb.2017.0345>
- Odegaard, B., & Lau, H. (2016). Methodological Considerations to Strengthen Studies of Peripheral Vision. *Trends in Cognitive Sciences*, 20(9), 642–643.
<http://doi.org/10.1016/j.tics.2016.06.005>
- Ohayon, M. M., & Schenck, C. H. (2010). Violent behavior during sleep: Prevalence, comorbidity and consequences. *Sleep Medicine*, 11(9), 941–946.
<http://doi.org/https://doi.org/10.1016/j.sleep.2010.02.016>
- Olivers, C. N., & Nieuwenhuis, S. (2005). The Beneficial Effect of Concurrent Task-Irrelevant Mental Activity on Temporal Attention. *Psychological Science*, 16(4), 265–269.
<http://doi.org/10.1111/j.0956-7976.2005.01526.x>
- Osho. (2003). *The ABC of Enlightenment: A Spiritual Dictionary for the Here and Now*. Harper Collins.
- Overgaard, M. (2010). How consciousness will change our view on neuroscience. *Cognitive Neuroscience*, 1(3), 224–225. <http://doi.org/10.1080/17588928.2010.497585>

- Overgaard, M., & Fazekas, P. (2016). Can No-Report Paradigms Extract True Correlates of Consciousness? *Trends in Cognitive Sciences*, 20(4), 241–242.
<http://doi.org/10.1016/j.tics.2016.01.004>
- Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, 15(4), 700–708.
<http://doi.org/10.1016/j.concog.2006.04.002>
- Pahnke, W. N. (1969). The Psychedelic Mystical Experience in the Human Encounter with Death. *The Harvard Theological Review*, 62(1), 1–21.
<http://doi.org/10.1017/S0017816000027577>
- Panayiotopoulos, C. P. (2008). Typical absence seizures and related epileptic syndromes: Assessment of current state and directions for future research. *Epilepsia*, 49(12), 2131–2139. <http://doi.org/10.1111/j.1528-1167.2008.01777.x>
- Parasuraman, R. (1998). The Attentive Brain: Issues and Prospects. In R. Parasuraman (Ed.), *The Attentive Brain* (pp. 3–15). MIT Press.
- Parfit, D. (1971). Personal identity. *Philosophical Review*, 80(1), 3–27.
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, 42(1), 107–123.
[http://doi.org/https://doi.org/10.1016/S0042-6989\(01\)00250-4](http://doi.org/https://doi.org/10.1016/S0042-6989(01)00250-4)
- Parnas, D. L. (2017). The Real Risks of Artificial Intelligence. *Communications of the ACM*, 60(10), 27–31. <http://doi.org/10.1145/3132724>
- Pashler, H. (1998). *The Psychology of Attention*. MIT Press.
- Pasquali, A., Timmermans, B., & Cleeremans, A. (2010). Know Thyself: Metacognitive Networks and Measures of Consciousness. *Cognition*, 117(2), 182–190.
<http://doi.org/doi:10.1016/j.cognition.2010.08.010>
- Pelli, D. G. (2008). Crowding: a cortical constraint on object recognition. *Current Opinion in Neurobiology*, 18(4), 445–451.
<http://doi.org/https://doi.org/10.1016/j.conb.2008.09.008>
- Perler, D. (2017). Self-Knowledge in Scholasticism. In U. Renz (Ed.), *Self-Knowledge: A History* (pp. 114–130). Oxford University Press.
<http://doi.org/10.1093/acprof:oso/9780190226411.003.0008>
- Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual Review of Neuroscience*, 35, 73–89. <http://doi.org/10.1146/annurev-neuro-062111-150525>
- Phillips, I. (2011a). Attention and Iconic Memory. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 204–227). Oxford University Press.
- Phillips, I. (2011b). Perception and Iconic Memory: What Sperling Doesn't Show. *Mind & Language*, 26(4), 381–411.
- Phillips, I. (2015). No watershed for overflow: Recent work on the richness of consciousness. *Philosophical Psychology*, 29(2), 1–14.
<http://doi.org/10.1080/09515089.2015.1079604>
- Phillips, I. (2018). The methodological puzzle of phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170347.
<http://doi.org/10.1098/rstb.2017.0347>
- Philosophy Now. (2012). Twenty-First Anniversary Survey. *Philosophy Now*, (93), 4–7.

- Retrieved from https://philosophynow.org/issues/93/Twenty-First_Anniversary_Survey
- Pierson, L. M., & Trout, M. (2017). What is consciousness for? *New Ideas in Psychology*, 47, 62–71. <http://doi.org/10.1016/j.newideapsych.2017.05.004>
- Pillai, M. (2016). Memory Athletes: Total Recall. Retrieved April 28, 2017, from <http://www.openthemagazine.com/article/sports/memory-athletes-total-recall>
- Pincham, H. L., Bowman, H., & Szucs, D. (2016). The experiential blink: Mapping the cost of working memory encoding onto conscious perception in the attentional blink. *Cortex*, 81, 35–49. <http://doi.org/10.1016/j.cortex.2016.04.007>
- Pinto, Y., Vandenbroucke, A. R., Otten, M., Sligte, I. G., Seth, A. K., & Lamme, V. A. (2017). Conscious visual memory with minimal attention. *Journal of Experimental Psychology: General*, 146(2), 214–226.
- Pitts, M. A., Lutsyshyna, L. A., & Hillyard, S. A. (2018). The relationship between attention and consciousness: an expanded taxonomy and implications for ‘no-report’ paradigms. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170348. <http://doi.org/10.1098/rstb.2017.0348>
- Pitts, M. A., Metzler, S., & Hillyard, S. A. (2014). Isolating neural correlates of conscious perception from neural correlates of reporting one’s perception. *Frontiers in Psychology*, 5(1078), 1–16. <http://doi.org/10.3389/fpsyg.2014.01078>
- Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S. A. (2014). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, 101, 337–350. <http://doi.org/https://doi.org/10.1016/j.neuroimage.2014.07.024>
- Poldrack, R. A., & Yarkoni, T. (2016). From Brain Maps to Cognitive Ontologies: Informatics and the Search for Mental Structure. *Annual Review of Psychology*, 67, 587–612.
- Poletti, M., Rucci, M., & Carrasco, M. (2017). Selective attention within the foveola. *Nature Neuroscience*, 20(10), 1413–1417. <http://doi.org/10.1038/nn.4622>
- Polger, T. W. (2017). Rethinking the Evolution of Consciousness. In S. Schneider & M. Velmans (Eds.), *The Blackwell Companion to Consciousness* (pp. 77–92). Wiley-Blackwell.
- Posner, M. I. (1994). Attention: The Mechanisms of Consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, 91(16), 7398–7403. Retrieved from <http://www.jstor.org/stable/2365291>
- Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience*, 13, 25–42.
- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, 139(1), 23–38. <http://doi.org/10.1016/j.neuroscience.2005.06.005>
- Powledge, T. M. (1997). Unlocking the secrets of the brain. *Bioscience*, 47(6), 330–334.
- Prinz, J. (2011). Is Attention Necessary and Sufficient for Consciousness? In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 174–203). Oxford University Press.
- Prinz, J. (2012). *The Conscious Brain: How Attention Engenders Experience*. Oxford University Press.
- Rabagliati, H., Robertson, A., & Carmel, D. (2018). The Importance of Awareness for Understanding Language. *Journal of Experimental Psychology: General*, 147(2), 190–208. <http://doi.org/10.1037/xge0000348>

- Raffman, D. (1993). *Language, Music, and Mind*. MIT Press.
- Ram-Prasad, C. (2007). *Indian Philosophy and the Consequences of Knowledge: Themes in Ethics, Metaphysics and Soteriology*. Routledge.
- Ramachandran, V. S., & Hirstein, W. (1997). Three laws of qualia: What neurology tells us about the biological functions of consciousness. *Journal of Consciousness Studies*, 4(5–6), 429–457.
- Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge University Press.
- Ramsøy, T., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, 3(1), 1–23.
<http://doi.org/10.1023/B:PHEN.0000041900.30172.e8>
- Rausch, M., Müller, H. J., & Zehetleitner, M. (2015). Metacognitive sensitivity of subjective reports of decisional confidence and visual experience. *Consciousness and Cognition*, 35, 192–205. <http://doi.org/10.1016/j.concog.2015.02.011>
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary Suppression of Visual Processing in an RSVP Task: An Attentional Blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849–860.
- Raymont, P., & Brook, A. (2009). Unity of Consciousness. In B. McLaughlin, A. Beckermann, & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 565–577). Oxford University Press.
- Rees, G. (1999). The Attentive Brain. *Trends in Cognitive Sciences*, 3(1), 40.
- Regan, M. A., Hallett, C., & Gordon, C. P. (2011). Driver distraction and driver inattention: Definition, relationship and taxonomy. *Accident Analysis and Prevention*, 43(5), 1771–1781. <http://doi.org/10.1016/j.aap.2011.04.008>
- Reiss, J., & Sprenger, J. (2017). Scientific Objectivity. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/win2017/entries/scientific-objectivity/>
- Reith, F., Brande, R., Synnot, A., Gruen, R., & Maas, A. (2016). The reliability of the Glasgow Coma Scale: a systematic review. *Intensive Care Medicine*, 42(1), 3–15.
<http://doi.org/DOI 10.1007/s00134-015-4124-3>
- Rensink, R. A. (2015). A Function-Centered Taxonomy of Visual Attention. In P. Coates & S. Coleman (Eds.), *Phenomenal Qualities: Sense, Perception, and Consciousness* (pp. 347–374). Oxford University Press. <http://doi.org/10.1093/acprof>
- Repovš, G., & Bresjanac, M. (2006). Cognitive neuroscience of working memory: A prologue. *Neuroscience*, 139(1), 1–3. <http://doi.org/10.1016/j.neuroscience.2005.12.007>
- Rescorla, M. (2017). The Computational Theory of Mind. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2017/entries/computational-mind/>
- Rhodes, S., & Cowan, N. (2018). Attention in working memory: attention is needed but it yearns to be free. *Annals of the New York Academy of Sciences*, 1424, 52–63.
<http://doi.org/10.1111/nyas.13652>
- Ribot, T. H. (1890). *The Psychology of Attention*. Open Court Publishing.
- Ric, F., & Muller, D. (2012). Unconscious Addition: When We Unconsciously Initiate and Follow Arithmetic Rules. *Journal Of Experimental Psychology-General*, 141(2), 222–226.

- Richards, B. (2013). Identity-Crowding and Object-Seeing: A Reply to Block. *Thought: A Journal of Philosophy*, 2(2), 9–19. <http://doi.org/10.1002/tht.57>
- Richards, B. (2015). Advancing the Overflow Debate. *Journal of Consciousness Studies*, 22(7–8), 124–144.
- Richards, B. (2016). Attention and seeing objects: The identity-crowding debate. *Philosophical Psychology*, 1–16. <http://doi.org/10.1080/09515089.2016.1150449>
- Richardson, J. TE. (2007). Measures of short-term memory: A historical review. *Cortex*, 43(5), 635–650.
- Ricker, T. J., Vergauwe, E., & Cowan, N. (2016). Decay Theory of Immediate Memory: From Brown (1958) to Today (2014). *Quarterly Journal of Experimental Psychology*, 69(10), 1969–1995. <http://doi.org/10.1080/17470218.2014.914546>
- Robinson, Z., Maley, C. J., & Piccinini, G. (2015). Is Consciousness a Spandrel? *Journal of the American Philosophical Association*, 1(2), 365–383.
- Roediger, H. L. (2008). Relativity of Remembering: Why the Laws of Memory Vanished. *Annual Review of Psychology*, 59, 225–254.
- Roediger, H., Rajaram, S., & Geraci, L. (2007). Three Forms of Consciousness in Retrieving Memories. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 251–288). Cambridge University Press.
- Rose, D. (2006). *Consciousness: Philosophical, Psychological and Neural Theories*. Oxford University Press.
- Rosenberg, M., Finn, E., Scheinost, D., Constable, R., & Chun, M. (2017). Characterizing Attention with Predictive Network Models. *Trends in Cognitive Sciences*, 21(4), 290–302.
- Rosenthal, D. M. (2009). Higher-order Theories of Consciousness. In B. P. McLaughlin, A. Beckermann, & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 239–252). Oxford University Press.
- Rosenthal, D. M. (2011). Exaggerated reports: reply to Block. *Analysis*, 71(3), 431–437. <http://doi.org/10.1093/analys/anr039>
- Rossi, E. L. (1998). The Feigenbaum scenario as a model of the limits of conscious information processing. *Bio Systems*, 46(1–2), 113–22. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9648682>
- Rovane, C. (2012). Does Rationality Enforce Identity? In A. Coliva (Ed.), *The Self and Self-Knowledge* (pp. 17–38). Oxford: Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199590650.003.0002>
- Rowlands, M. (2015). Rilkean Memory. *Southern Journal of Philosophy*, 53(S1), 141–154. <http://doi.org/10.1111/sjp.12118>
- Ruff, C. C. (2011). A Systems-Neuroscience View of Attention. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 1–23). Oxford University Press.
- Saiki, J., & Holcombe, A. O. (2012). Blindness to a simultaneous change of all elements in a scene, unless there is a change in summary statistics. *Journal of Vision*, 12(3), 2:1–11. Retrieved from <http://www.journalofvision.org/content/12/3/2>
- Samaha, J., Barrett, J. J., Sheldon, A. D., LaRocque, J. J., & Postle, B. R. (2016). Dissociating perceptual confidence from discrimination accuracy reveals no influence of metacognitive awareness on working memory. *Frontiers in Psychology*, 7(851), 1–8. <http://doi.org/10.3389/fpsyg.2016.00851>

- Sawaki, R., & Luck, S. (2010). Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal. *Attention, Perception, & Psychophysics*, 72(6), 1455–1470. <http://doi.org/10.3758/APP.72.6.1455>
- Schelonka, K., Graulty, C., Canseco-Gonzalez, E., & Pitts, M. A. (2017). ERP signatures of conscious and unconscious word and letter perception in an inattentional blindness paradigm. *Consciousness and Cognition*, 54, 56–71.
- Schlagbauer, B., Rausch, M., Zehetleitner, M., Müller, H. J., & Geyer, T. (2018). Contextual cueing of visual search is associated with greater subjective experience of the search display configuration. *Neuroscience of Consciousness*, 2018(1), niy001. <http://doi.org/10.1093/nc/niy001>
- Schmitz, T. W., & Duncan, J. (2018). Normalization and the Cholinergic Microcircuit: A Unified Basis for Attention. *Trends in Cognitive Sciences*, 22(5), 422–437. <http://doi.org/10.1016/j.tics.2018.02.011>
- Schooler, J. W. (2002). Re-representing consciousness: dissociations between experience and meta-consciousness. *Trends in Cognitive Sciences*, 6(8), 339–344. [http://doi.org/10.1016/S1364-6613\(02\)01949-6](http://doi.org/10.1016/S1364-6613(02)01949-6)
- Schroer, R. (2012). Representationalism and the scene-immediacy of visual experience: A journey to the fringe and back. *Philosophical Psychology*, 25(4), 595–615. <http://doi.org/10.1080/09515089.2011.625120>
- Schwartz, S. H. (2017). *Visual Perception: A Clinical Orientation* (5th ed.). McGraw-Hill.
- Schwitzgebel, E. (2007). Do You Have Constant Tactile Experience of Your Feet in Your Shoes? Or Is Experience Limited to What's in Attention? *Journal of Consciousness Studies*, 14(3), 5–35.
- Schwitzgebel, E. (2011). *Perplexities of consciousness*. Cambridge, Mass.: MIT Press.
- Scimeca, J. M., Kiyonaga, A., & D'Esposito, M. (2018). Reaffirming the Sensory Recruitment Account of Working Memory. *Trends in Cognitive Sciences*, 22(3), 190–192. <http://doi.org/10.1016/j.tics.2017.12.007>
- Seager, W. (2007). A Brief History of the Philosophical Problem of Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 9–34). Cambridge University Press.
- Seager, W. (2015). Montemayor, Carlos. Consciousness, Attention, and Conscious Attention. *CHOICE: Current Reviews for Academic Libraries*, 53(4), 649.
- Seager, W. (2016). *Theories of Consciousness : An Introduction and Assessment*. London, UNITED KINGDOM: Routledge. Retrieved from <http://ebookcentral.proquest.com/lib/mqu/detail.action?docID=4406186>
- Searle, J. R. (1992). *The rediscovery of the mind*. MIT Press.
- Searle, J. R. (1993). The Problem of Consciousness. *Consciousness and Cognition*, 2(4), 310–319. <http://doi.org/10.1006/ccog.1993.1026>
- Seli, P., Kane, M. J., Smallwood, J., Schacter, D. L., Maillet, D., Schooler, J. W., & Smilek, D. (2018). Mind-Wandering as a Natural Kind: A Family-Resemblances View. *Trends in Cognitive Sciences*, 22(6), 479–490. <http://doi.org/10.1016/j.tics.2018.03.010>
- Seneviratne, U., Woo, J. J., Boston, R. C., Cook, M., & D'souza, W. (2015). Focal seizure symptoms in idiopathic generalized epilepsies. *Neurology*, 85(7), 589–595.
- Serences, J. T. (2016). Neural mechanisms of information storage in visual short-term memory. *Vision Research*, 128, 53–67. <http://doi.org/https://doi.org/10.1016/j.visres.2016.09.010>

- Serences, J. T., & Kastner, S. (2014). A Multi-level Account of Selective Attention. In A. Nobre & S. Kastner (Eds.), *The Oxford Handbook of Attention* (pp. 76–104). Oxford University Press.
- Sergent, C., & Dehaene, S. (2004). Is Consciousness a Gradual Phenomenon? Evidence for an All-or-None Bifurcation During the Attentional Blink. *Psychological Science*, 15(11), 720–728. <http://doi.org/10.1111/j.0956-7976.2004.00748.x>
- Sergent, C., & Rees, G. (2007). Conscious access overflows overt report. *Behavioral and Brain Sciences*, 30(5–6), 523–524. <http://doi.org/10.1017/S0140525X07003044>
- Service, E. (1998). The Effect of Word Length on Immediate Serial Recall Depends on Phonological Complexity, Not Articulatory Duration. *The Quarterly Journal of Experimental Psychology Section A*, 51(2), 283–304. <http://doi.org/10.1080/713755759>
- Shanahan, M. (2016). The Frame Problem. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 201). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/spr2016/entries/frame-problem/>
- Shanon, B. (2002). *The antipodes of the mind: charting the phenomenology of the Ayahuasca experience*. Oxford University Press.
- Shea, N., & Bayne, T. (2010). The vegetative state and the science of consciousness. *The British Journal for the Philosophy of Science*, 61(3), 459–484. <http://doi.org/10.1093/bjps/axp046>
- Shoemaker, S. (1994). Self-knowledge and "inner sense": Lecture I: The object perception model. *Philosophy and Phenomenological Research*, 54(2), 249–269.
- Siclari, F., LaRocque, J., Postle, B., & Tononi, G. (2013). Assessing sleep consciousness within subjects using a serial awakening paradigm. *Frontiers in Psychology*, 4, 542. <http://doi.org/10.3389/fpsyg.2013.00542>
- Siewert, C. (2012). On the Phenomenology of Introspection. In D. Smithies & D. Stoljar (Eds.), *Introspection and Consciousness*. Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199744794.003.0005>
- Silins, N. (2015). Perceptual Experience and Perceptual Justification. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 201). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/win2015/entries/perception-justification/>
- Simon, H. A. (1980). Cognitive Science: The Newest Science of the Artificial. *Cognitive Science*, 4(1), 33–46. [http://doi.org/10.1016/S0364-0213\(81\)80003-1](http://doi.org/10.1016/S0364-0213(81)80003-1)
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, 28(9), 1059–1074. <http://doi.org/10.1068/p281059>
- Simons, D. J., Hannula, D. E., Warren, D. E., & Day, S. W. (2007). Behavioral, Neuroimaging, and Neuropsychological Approaches to Implicit Perception. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge Handbook of Consciousness* (pp. 207–250). Cambridge University Press.
- Singer, P. (1975). *Animal liberation: A new ethics for our treatment of animals*. Random House.
- Singleton, A., Mason, M., & Webber, R. (2004). Spirituality in adolescence and young adulthood: A method for a qualitative study. *International Journal of Children's Spirituality*, 9(3), 247–262.
- 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>
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLOS ONE*, 3(2), e1699.
- Sligte, I. G., Vandenbroucke, A., Scholte, H. S., & Lamme, V. (2010). Detailed Sensory Memory, Sloppy Working Memory. *Frontiers in Psychology*, 1(175), 1–10. <http://doi.org/10.3389/fpsyg.2010.00175>
- Smart, J. (1959). Can biology be an exact science? *Synthese*, 11(4), 359–368.
- Smith, N. J. (2012). *Logic: The Laws of Truth*. Princeton University Press.
- Smithies, D. (2011). Attention Is Rational-Access Consciousness. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 247–273). Oxford University Press.
- Smithies, D., & Stoljar, D. (Eds.). (2012). *Introspection and Consciousness*. Oxford University Press.
- Soteriou, M. (2013). *The mind's construction: the ontology of mind and mental action*. Oxford University Press.
- Soto, D., Mäntylä, T., & Silvanto, J. (2011). Working memory without consciousness. *Current Biology*, 21(22), R912–R913. <http://doi.org/10.1016/j.cub.2011.09.049>
- Soto, D., & Silvanto, J. (2014). Reappraising the relationship between working memory and conscious awareness. *Trends in Cognitive Sciences*, 18(10), 520–525. <http://doi.org/10.1016/j.tics.2014.06.005>
- Soto, D., & Silvanto, J. (2016). Is conscious awareness needed for all working memory processes? *Neuroscience of Consciousness*, 2016(1), 1–3. <http://doi.org/10.1093/nc/niw009>
- Souza, A. S., Rerko, L., & Oberauer, K. (2014). Unloading and reloading working memory: Attending to one item frees capacity. *Journal of Experimental Psychology: Human Perception and Performance*, 40(3), 1237–1256. <http://doi.org/10.1037/a0036331>
- Sperling, G. (1960). The Information Available in Brief Visual Presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29.
- Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, 82(3), 171–177.
- Squire, L. R. (2009). Memory and Brain Systems: 1969–2009. *The Journal of Neuroscience*, 29(41), 12711–12716. <http://doi.org/10.1523/JNEUROSCI.3575-09.2009>
- Staugaard, C. F., Petersen, A., & Vangkilde, S. (2016). Eccentricity effects in vision and attention. *Neuropsychologia*, 92, 69–78. <http://doi.org/https://doi.org/10.1016/j.neuropsychologia.2016.06.020>
- Stazicker, D. J. (2011a). Attention, Visual Consciousness and Indeterminacy. *Mind and Language*, 26(2), 156–184. <http://doi.org/10.1111/j.1468-0017.2011.01414.x>
- Stazicker, D. J. (2011b). *Attention and the Indeterminacy of Visual Experience*. University of California, Berkeley.
- Stazicker, D. J. (2018). Partial report is the wrong paradigm. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170350. <http://doi.org/10.1098/rstb.2017.0350>
- Stoljar, D. (2016). The Semantics of ‘What it’s like’ and the Nature of Consciousness. *Mind*, 125(500), 1161–1198.

- Storm, J. F., Boly, M., Casali, A. G., Massimini, M., Olcese, U., Pennartz, C. M. A., & Wilke, M. (2017). Consciousness regained: Disentangling mechanisms, brain systems, and behavioral responses, 37(45). <http://doi.org/10.1523/JNEUROSCI.1838-17.2017>
- Strawson, G. (1994). *Mental Reality*. MIT Press.
- Strawson, G. (2016). Mind and Being—the Primacy of Panpsychism. In Brüntrup & Jaskolla (Eds.), *Panpsychism: Contemporary Perspectives*. Oxford University Press. <http://doi.org/10.1093/acprof:oso/9780199359943.003.0004>
- Strawson, G. (2017). Physicalist Panpsychism. In S. Schneider & M. Velmans (Eds.), *The Blackwell Companion to Consciousness* (2nd ed., pp. 374–390). Blackwell.
- Styles, E. A. (1997). *The Psychology of Attention*. Psychology Press.
- Sumner, P., Tsai, P.-C., Yu, K., & Nachev, P. (2006). Attentional modulation of sensorimotor processes in the absence of perceptual awareness. *Proceedings of the National Academy of Sciences*, 103(27), 10520–10525. <http://doi.org/10.1073/pnas.0601974103>
- Swancer, B. (2017, July 17). Miraculous Cases of People Who Lived Without a Brain. Retrieved October 10, 2018, from <https://mysteriousuniverse.org/2017/07/miraculous-cases-of-people-who-lived-without-a-brain/>
- Syková, E., & Vargová, L. (2008). Extrasynaptic transmission and the diffusion parameters of the extracellular space. *Neurochemistry International*, 52(1–2), 5–13. <http://doi.org/10.1016/j.neuint.2007.04.007>
- Tallon-Baudry, C. (2012). On the neural mechanisms subserving consciousness and attention. *Frontiers in Psychology*, 2(397), 1–11. <http://doi.org/10.3389/fpsyg.2011.00397>
- Taylor, H. (2018). Attention, Psychology, and Pluralism. *The British Journal for the Philosophy of Science*, 69(4), 935–956. Retrieved from <http://dx.doi.org/10.1093/bjps/axx030>
- Taylor, J. G. (2002). Paying attention to consciousness. *Trends in Cognitive Sciences*, 6(5), 206–210.
- Taylor, J. G. (2003). Paying attention to consciousness. *Progress in Neurobiology*, 71(4), 305–335. <http://doi.org/10.1016/j.pneurobio.2003.10.002>
- Taylor, J. H. (2013a). Is Attention Necessary and Sufficient for Phenomenal Consciousness? *Journal of Consciousness Studies*, 20(11–12), 173–194.
- Taylor, J. H. (2013b). Is the grain of vision finer than the grain of attention? Response to Block. *Thought: A Journal of Philosophy*, 2, 20–28. <http://doi.org/10.1002/tht.59>
- Taylor, J. H. (2015a). Against Unifying Accounts of Attention. *Erkenntnis*, 80, 39–56. <http://doi.org/10.1007/s10670-014-9611-3>
- Taylor, J. H. (2015b). *Consciousness and the Heterogeneity of Attention*. Durham University. Retrieved from <http://etheses.dur.ac.uk/10973/>
- Tellegen, A., & Waller, N. (2008). Exploring personality through test construction: Development of the Multidimensional Personality Questionnaire. In G. Boyle, G. Matthews, & D. Saklofske (Eds.), *The Sage handbook of personality theory and assessment: Personality measurement and testing, Volume 2* (pp. 261–292). SAGE Publications.
- Theeuwes, J. (1991). Exogenous and endogenous control of attention: the effect of visual onsets and offsets. *Perception & Psychophysics*, 49(1), 83–90.

- Theeuwes, J., Olivers, C. N., & Chizk, C. L. (2005). Remembering a Location Makes the Eyes Curve Away. *Psychological Science*, 16(3), 196–199. <http://doi.org/10.1111/j.0956-7976.2005.00803.x>
- Thompson, E. (2015a). Dreamless Sleep, the Embodied Mind, and Consciousness. In T. K. Metzinger & J. M. Windt (Eds.), *Open MIND* (pp. 1–19). Frankfurt am Main: MIND Group. <http://doi.org/10.15502/9783958570351>
- Thompson, E. (2015b). *Waking, dreaming, being: Self and consciousness in neuroscience, meditation, and philosophy*. Columbia University Press.
- Thompson, V. A., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology*, 63(3), 107–140. <http://doi.org/10.1016/j.cogpsych.2011.06.001>
- Tiffin, H. (2016). Do Insects Feel Pain? *Animal Studies Journal*, 5(1), 80–96. Retrieved from <http://ro.uow.edu.au/asj/vol5/iss1/6>
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208. <http://doi.org/10.1037/h0061626>
- Tong, F. (2013). Imagery and visual working memory: one and the same? *Trends in Cognitive Sciences*, 17(10), 489–490. <http://doi.org/10.1016/j.tics.2013.08.005>
- Tononi, G. (2008). Consciousness as Integrated Information: a Provisional Manifesto. *The Biological Bulletin*, 215(3), 216–242. <http://doi.org/10.2307/25470707>
- Tononi, G., & Laureys, S. (2009). The Neurology of Consciousness: An Overview. In S. Laureys & G. Tononi (Eds.), *The Neurology of Consciousness: Cognitive Neuroscience and Neuropathology* (pp. 375–412). Elsevier.
- Travis, F., & Shear, J. (2010). Reply to Josipovic: Duality and non-duality in meditation research. *Consciousness and Cognition*, 19(4), 1122–1123.
- Travis, L. (1937). Brain potentials and the temporal course of consciousness. *Journal of Experimental Psychology*, 21(3), 302–309. <http://doi.org/10.1037/h0056742>
- Treffert, D., & Christensen, D. (2005). Inside the Mind of a Savant. *Scientific American*, 293(6), 108–113.
- Treisman, A. M. (1964). Selective Attention in Man. *British Medical Bulletin*, 20(1), 12–16.
- Treisman, A. M. (1969). Strategies and Models of Selective Attention. *Psychological Review*, 76(3), 282–299.
- Treisman, A. M. (2003). Consciousness and Perceptual Binding. In A. Cleeremans (Ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. (pp. 95–113). Oxford University Press.
- Treisman, A. M., & Gelade, G. (1980). A Feature-Integration Theory of Attention. *Cognitive Psychology*, 12, 97–136.
- Trevarthen, C., & Reddy, V. (2017). Consciousness in Infants. In S. Schneider & M. Velmans (Eds.), *The Blackwell Companion to Consciousness* (2nd ed., pp. 45–62). Wiley-Blackwell.
- Trübutschek, D., Marti, S., Ojeda, A., King, J.-R., Mi, Y., Tsodyks, M., & Dehaene, S. (2017). A theory of working memory without consciousness or sustained activity. *ELife*, 6(e23871), 1–29. <http://doi.org/10.7554/eLife.23871>
- Tsuchiya, N., & van Boxtel, J. J. (2010). Is recurrent processing necessary and/or sufficient for consciousness? *Cognitive Neuroscience*, 1(3), 230–231. <http://doi.org/10.1080/17588928.2010.497582>

- Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. (2015). No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences*, 19(12), 757–770. <http://doi.org/10.1016/j.tics.2015.10.002>
- Tulving, E. (1985). How Many Memory Systems Are There ? *American Psychologist*, 40(4), 385–398.
- Tulving, E. (2007). Are There 256 Different Kinds of Memory? In J. S. Nairne (Ed.), *The Foundations of Remembering: Essays in Honor of Henry L. Roediger III* (pp. 39–52). Psychology Press.
- Turausky, K. E. (2014). The Most Interesting Problem in the Universe. *Journal of Consciousness Studies*, 21(7–8), 220–240.
- Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. MIT Press.
- Tye, M. (2009). A new look at the speckled hen. *Analysis*, 69(2), 258–263.
- Tye, M. (2010). Attention, Seeing, and Change Blindness. *Philosophical Issues*, 20(1), 410–437. <http://doi.org/10.1111/j.1533-6077.2010.00180.x>
- Tye, M. (2014). Does Conscious Seeing Have A Finer Grain Than Attention? *Thought: A Journal of Philosophy*, 3(2), 154–158. <http://doi.org/10.1002/tht3.123>
- Usher, M., Bronfman, Z. Z., Talmor, S., Jacobson, H., & Eitam, B. (2018). Consciousness without report: insights from summary statistics and inattention “blindness.” *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170354. <http://doi.org/10.1098/rstb.2017.0354>
- van Boxtel, J. J., Tsuchiya, N., & Koch, C. (2010). Consciousness and attention: On sufficiency and necessity. *Frontiers in Psychology*, 1, 1–13. <http://doi.org/10.3389/fpsyg.2010.00217>
- van Es, D. M., Theeuwes, J., Knapen, T., & Ré, D. A. (2018). Spatial sampling in human visual cortex is modulated by both spatial and feature-based attention. *eLife*, 7, 1–28. <http://doi.org/10.7554/eLife.36928.001>
- van Gaal, S., De Lange, F., & Cohen, M. (2012). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*, 6(121), 1–15. <http://doi.org/10.3389/fnhum.2012.00121>
- van Gulick, R. (2017). Consciousness. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Summer). Stanford University. Retrieved from <https://plato.stanford.edu/archives/sum2017/entries/consciousness/>
- van Swinderen, B. (2005). The remote roots of consciousness in fruit-fly selective attention? *BioEssays*, 27(3), 321–330. <http://doi.org/10.1016/B978-044452977-0/50003-6>
- Vandenbroucke, A. R., Sligte, I. G., Barrett, A. B., Seth, A. K., Fahrenfort, J. J., & Lamme, V. A. (2014). Accurate metacognition for visual sensory memory representations. *Psychological Science*, 25(4), 861–873. <http://doi.org/10.1177/0956797613516146>
- Vandenbroucke, A. R., Sligte, I. G., & Lamme, V. A. (2011). Manipulations of attention dissociate fragile visual short-term memory from visual working memory. *Neuropsychologia*, 49(6), 1559–1568. <http://doi.org/10.1016/j.neuropsychologia.2010.12.044>
- Velmans, M. (2012). The evolution of consciousness. *Contemporary Social Science*, 7(2), 117–138. <http://doi.org/10.1080/21582041.2012.692099>
- Velmans, M. (2014). Conscious Agency and the Preconscious/Unconscious Self. In S. Menon, A. Sinha, & B. V. Sreekantan (Eds.), *Interdisciplinary Perspectives on Consciousness and the Self* (pp. 11–25). Springer India. Retrieved from

- <http://link.springer.com.simsrad.net.ocs.mq.edu.au/book/10.1007%2F978-81-322-1587-5>
- Vliegenthart, D. (2011). Can Neurotheology Explain Religion? *Archive for the Psychology of Religion*, 33(2), 137–171. <http://doi.org/10.1163/157361211X577284>
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of Features, Conjunctions, and Objects in Visual Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27(1), 92–114. <http://doi.org/10.1037/0096-1523.27.1.92>
- Wackermann, J., Pütz, P., & Allefeld, C. (2008). Ganzfeld-induced hallucinatory experience, its phenomenology and cerebral electrophysiology. *Cortex*, 44(10), 1364–1378. <http://doi.org/10.1016/j.cortex.2007.05.003>
- Wager, T. D., Jonides, J., Smith, E. E., & Nichols, T. E. (2005). Toward a taxonomy of attention shifting: Individual differences in fMRI during multiple shift types. *Cognitive, Affective & Behavioral Neuroscience*, 5(2), 127–43. <http://doi.org/10.3758/CABN.5.2.127>
- Walther, S., & Strik, W. (2016). Catatonia. *CNS Spectrums*, 21(4), 341–348. <http://doi.org/10.1017/S1092852916000274>
- Ward, E. J. (2018). Downgraded phenomenology: how conscious overflow lost its richness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170355. <http://doi.org/10.1098/rstb.2017.0355>
- Wassell, J., Rogers, S. L., Felmingam, K. L., Bryant, R. A., & Pearson, J. (2015). Sex hormones predict the sensory strength and vividness of mental imagery. *Biological Psychology*, 107(C), 61–68. <http://doi.org/10.1016/j.biopsycho.2015.02.003>
- Watts, A. (1964). A psychedelic experience: fact or fantasy? In D. Solomon (Ed.), *LSD: The Consciousness-Expanding Drug*. Putnam.
- Watzl, S. (2010). *The Significance of Attention*. Columbia University.
- Watzl, S. (2011a). Attention as Structuring of the Stream of Consciousness. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 145–173). Oxford University Press.
- Watzl, S. (2011b). The Nature of Attention. *Philosophy Compass*, 6(11), 842–853. <http://doi.org/10.1111/j.1747-9991.2011.00433.x>
- Wegner, D. M., & Smart, L. (1997). Deep cognitive activation: a new approach to the unconscious. *Journal of Consulting and Clinical Psychology*, 65(6), 984–995.
- Weldon, K. B., Rich, A. N., Woolgar, A., & Williams, M. A. (2016). Disruption of Foveal Space Impairs Discrimination of Peripheral Objects. *Frontiers in Psychology*, 7, 699. <http://doi.org/10.3389/fpsyg.2016.00699>
- Wertheim, M. (2016). I feel therefore I am. Retrieved August 8, 2016, from <https://aeon.co/essays/how-and-why-exactly-did-consciousness-become-a-problem>
- Whitney, D., Haberman, J., & Sweeny, T. (2014). From textures to crowds: multiple levels of summary statistical perception. In J. Werner & L. Chalupa (Eds.), *The New Visual Neurosciences* (pp. 695–710). MIT Press.
- 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/https://doi.org/10.1016/j.tics.2011.02.005>
- Whitney, D., & Yamanashi Leib, A. (2018). Ensemble Perception. *Annual Review of Psychology*, 69(1), 105–129. <http://doi.org/10.1146/annurev-psych-010416-044232>

- Wider, K. (1997). *The Bodily Basis of Consciousness*. Cornell University Press.
- Wilke, M., Mueller, K.-M., & Leopold, D. A. (2009). Neural activity in the visual thalamus reflects perceptual suppression. *Proceedings of the National Academy of Sciences*, 106(23), 9465–9470.
- Wilkes, K. (1984). Is Consciousness Important? *British Journal for the Philosophy of Science*, 35(3), 223–243. <http://doi.org/10.1093/bjps/35.3.223>
- Windey, B., & Cleeremans, A. (2015). Consciousness as a graded and an all-or-none phenomenon: A conceptual analysis. *Consciousness and Cognition*, 35, 185–191. <http://doi.org/10.1016/j.concog.2015.03.002>
- Windt, J. M. (2015). Just in Time—Dreamless Sleep Experience as Pure Subjective Temporality. In T. K. Metzinger & J. M. Windt (Eds.), *Open MIND* (pp. 1–34). Frankfurt am Main: MIND Group. <http://doi.org/10.15502/9783958571174>
- Windt, J. M., Nielsen, T., & Thompson, E. (2016). Does Consciousness Disappear in Dreamless Sleep? *Trends in Cognitive Sciences*, 20(12), 871–882. <http://doi.org/https://doi.org/10.1016/j.tics.2016.09.006>
- Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, 1(2), 202–238.
- Wolfe, J. M. (1999a). Inattentional Amnesia. In V. Coltheart (Ed.), *Fleeting Memories* (pp. 71–94). MIT Press.
- Wolfe, J. M. (1999b). Visual experience: Less than you think, more than you remember. In C. Taddei-Feretti & C. Musio (Eds.), *Neuronal basis and psychological aspects of consciousness* (pp. 165–185). World Scientific.
- Woodman, G. F., & Luck, S. J. (2003). Dissociations Among Attention, Perception, and Awareness During Object-Substitution Masking. *Psychological Science*, 14(6), 605–611. <http://doi.org/10.1046/j.0956-7976.2003.psci>
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press. <http://doi.org/10.1017/CBO9781107415324.004>
- Woodward, J. (2014). Scientific Explanation. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Winter). Stanford University. <http://doi.org/10.1111/1467-9973.00225>
- Wu, C.-C., & Wolfe, J. M. (2018). A New Multiple Object Awareness Paradigm Shows that Imperfect Knowledge of Object Location Is Still Knowledge. *Current Biology*, 28(21), 3430–3434.e3. <http://doi.org/10.1016/j.cub.2018.08.042>
- Wu, W. (2011). Attention as Selection for Action. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays* (pp. 97–116). Oxford University Press.
- Wu, W. (2013). The Conscious Brain: How Attention Engenders Experience, by Jesse Prinz. *Mind*, 122(488), 1174–1180. <http://doi.org/10.1093/mind/fzt101>
- Wu, W. (2014). *Attention*. Routledge.
- Wu, W. (2016). *A Primer on Attention as Selection for Action*. Retrieved from <https://www.academia.edu/s/f6760af11b>
- Wundt, W. M. (1912). *An introduction to psychology*. (R. Pintner, Ed.). New York, N.Y.: MacMillan Co.
- Wyart, V., Dehaene, S., & Tallon-Baudry, C. (2012). Early dissociation between neural signatures of endogenous spatial attention and perceptual awareness during visual masking. *Frontiers in Human Neuroscience*, 6(16), 1–14. <http://doi.org/10.3389/fnhum.2012.00016>

- Wyart, V., & Tallon-Baudry, C. (2008). Neural Dissociation between Visual Awareness and Spatial Attention. *Journal of Neuroscience*, 28(10), 2667–2679.
<http://doi.org/10.1523/JNEUROSCI.4748-07.2008>
- Wyble, B. (2015). Presentation: Visual Attention (6 November). Perception Group, University of Sydney.
- Xu, Y. (2018). Sensory Cortex Is Nonessential in Working Memory Storage. *Trends in Cognitive Sciences*, 22(3), 192–193. <http://doi.org/10.1016/j.tics.2017.12.008>
- Yeshurun, Y. (2019). The spatial distribution of attention. *Current Opinion in Psychology*, 29, 76–81. <http://doi.org/10.1016/j.copsyc.2018.12.008>
- Ylikoski, P. (2013). Causal and Constitutive Explanation Compared. *Erkenntnis*, 78, 277–297.
- Yoshimi, J., & Vinson, D. W. (2015). Extending Gurwitsch's field theory of consciousness. *Consciousness and Cognition*, 34, 104–123.
<http://doi.org/10.1016/j.concog.2015.03.017>
- Zahavi, D. (2008). Simulation, projection and empathy. *Consciousness and Cognition*, 17(2), 514–522. <http://doi.org/https://doi.org/10.1016/j.concog.2008.03.010>
- Zelazo, P. D. (2004). The development of conscious control in childhood. *Trends in Cognitive Sciences*, 8(1), 12–17. <http://doi.org/10.1016/j.tics.2003.11.001>
- Zelazo, P. D., Gao, H. H., & Todd, R. (2007). The Development of Consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness* (pp. 405–432). Cambridge University Press.
- Zeman, A. (2001). Consciousness. *Brain: A Journal of Neurology*, 124(7), 1263–1289.
<http://doi.org/Doi 10.1093/Brain/124.7.1263>
- Zimmermann, J. F., Moscovitch, M., & Alain, C. (2016). Attending to auditory memory. *Brain Research*, 1640, 208–221. <http://doi.org/10.1016/j.brainres.2015.11.032>
- Zubek, J., Hughes, G., & Shephard, J. (1971). A comparison of the effects of prolonged sensory deprivation and perceptual deprivation. *Canadian Journal of Behavioural Science/Revue Canadienne Des Sciences Du Comportement*, 3(3), 282–290.
<http://doi.org/10.1037/h0082270>