



Annan, Heather May (2026) *Attentional differences in autism – the Priority Monotropism account*. PhD thesis.

<https://theses.gla.ac.uk/86260/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

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

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

Attentional Differences in Autism – the Priority Monotropism Account

Heather May Annan

MSc Philosophy of Mind and Psychology – University of Glasgow

BA Philosophy and Psychology – Open University

Submitted in fulfilment of the requirements for the degree of PhD Philosophy and
Psychology (Research)

School of Humanities

College of Arts and Humanities

University of Glasgow

July 2026

Abstract

Throughout this thesis, I will develop a novel essentialist theory of autism, Priority Monotropism, and demonstrate its explanatory power. I will first illustrate the necessity of a new theory of autism by demonstrating the harms that have resulted, and continue to result, from our misunderstanding of the condition.

I will introduce the heterogeneity problem and explain that it poses a serious threat to essentialist theories of autism. I will explain that the heterogeneity problem of autism has two prongs. The first prong of the heterogeneity problem is that autism is associated with a wide range of seemingly disparate symptoms and behaviours that do not share a clear conceptual link. The second prong of the heterogeneity problem is that there is individual variance between the symptom profiles of autistic individuals. I will introduce the following essentialist theories of autism: Weak Central Coherence theory, Theory of Mind theory, Enhanced Perceptual Functioning theory, and Executive Dysfunction theory. I will demonstrate that these theories are unable to overcome the heterogeneity problem. I will then introduce a further essentialist theory of autism—Monotropism. I will show that while Monotropism has the potential to overcome the heterogeneity problem, it is insufficient in its current form. I will claim that this is largely due to the lack of a clear definition of key terms, such as ‘attention’ and ‘interest’. This is especially problematic since this theory proposes that the core essence of autism is a difference in attentional style preference.

In order to resolve this issue, I will create a revised account of Monotropism that makes clearer claims about these terms. In order to do this, I will examine three philosophical theories of attention and assess their suitability for providing a foundation for a revised account of Monotropism. These three theories are: the Selection for Action theory, the Cognitive Unison theory, and the Priority Structure account. I will ultimately conclude that the Priority Structure account provides the best foundation, and I will develop a revised account of Monotropism—Priority Monotropism—by using this theory as a foundation. I will then demonstrate that Priority Monotropism can overcome both prongs of the heterogeneity problem. Additionally, I will also demonstrate that considering autism through the lens of Priority Monotropism can allow for potential explanations about other previously unexplained aspects of autism, such as why it commonly co-occurs with a wide range of seemingly unrelated conditions.

Table of Contents

Acknowledgements	11
Author's Declaration	12
Introduction	13
Terminology Disclaimer	16
Chapter 1 – Understanding Autism: Highlighting the Necessity of a New Theory of Autism	18
Section 1.1 – Chapter Introduction	18
Section 1.2 – Showing the Explanatory Gap	18
Section 1.3 – Theoretical Consequences	20
Section 1.3.1 – Autism Research Efforts	20
Section 1.3.2 – Eliminativism About Autism	20
Section 1.4 – Practical Consequences	23
Section 1.4.1 – Consequences for Autism Support	23
Section 1.4.1.1 – Employment	23
Section 1.4.1.2 – Education	25
Section 1.4.1.3 – Suicidal Ideation and Suicide Attempts	27
Section 1.4.1.4 – Victimization	29
Section 1.4.2 - Diagnosis	30
Section 1.4.2.1 – Current Diagnostic Criteria	30
Section 1.4.2.2 – Problems with the DSM-5 Criteria	31
Section 1.4.2.3 – Problems with Behavioural Observation	32
Section 1.4.2.3.1 – Gender	33
Section 1.4.2.3.2 – Ethnicity and Culture	35
Section 1.4.2.3.3 – Socio-economic Status	36
Section 1.4.2.3.4 – Age	36
Section 1.4.2.3.5 – Co-occurring Conditions	37
Section 1.5 – Harmful Misconceptions About Autism	39
Section 1.5.1 – General Misconceptions	39
Section 1.5.2 – Misconceptions Surrounding Vaccinations	40
Section 1.5.3 – Misconceptions About ‘Treatments’	43

Section 1.5.3.1 – Harmful Treatments: Chelation Therapy.....	43
Section 1.5.3.2 – Harmful Treatments: Secretin Injections.....	44
Section 1.5.3.3 – Harmful Treatments: Facilitated Communication.....	44
Section 1.5.3.4 – Harmful Treatments: Electroshock Therapy	45
Section 1.5.4 – Supernatural Ideology About Autism	46
Section 1.5.4.1 – Supernatural Ideology: Changelings	46
Section 1.5.4.2 – Supernatural Ideology: Demons and Possession.....	47
Section 1.5.5 – Misconceptions Based on Historic Autism Theories	47
Section 1.5.5.1 – Harmful Theories: Refrigerator Mother Theory	48
Section 1.5.5.2 – Harmful Theories: Extreme Male Brain Theory	49
Section 1.5.5.3 – Mistrust of Autism Research Efforts	50
Section 1.6 – Chapter Conclusion	50
 Chapter 2 – The Heterogeneity Problem of Autism: Essentialist Theories and Eliminativism	 51
Section 2.1 – Chapter Introduction.....	51
Section 2.2 – The Heterogeneity Problem	52
Section 2.2.1 – Listing the Wide and Varied Range of Symptoms and Behaviours in Autism.....	53
Section 2.2.1.1 – Symptoms and Behaviours Related to Differences in Communication	53
Section 2.2.1.2 – Symptoms and Behaviours Related to Differences in Behaviours and Interests.....	55
Section 2.2.1.3 – Symptoms and Behaviours Related to Sensory Differences	56
Section 2.2.1.4 – Symptoms and Behaviours Related to Stress	56
Section 2.2.1.5 – Symptoms And Behaviours Related to Difficulties with Fitting In	58
Section 2.2.1.6 – Symptoms And Behaviours Related to Differences in Motor Skills.....	58
Section 2.2.2 – The Heterogeneity Problem and the Difficulty of Defining Autism.....	59
Section 2.2.2.1 – The DSM-5 and the Cluster of Behaviours Model	59
Section 2.3 – Essentialist Theories of Autism	61
Section 2.3.1 – Weak Central Coherence Theory	62
Section 2.3.1.1 – Strengths of Weak Central Coherence Theory	64

Section 2.3.1.2 – Weaknesses of Weak Central Coherence Theory	64
Section 2.3.2 – Enhanced Perceptual Functioning Theory	66
Section 2.3.2.1 – Strengths of Enhanced Perceptual Functioning Theory	66
Section 2.3.2.2 – Weaknesses of Enhanced Perceptual Functioning Theory	66
Section 2.3.3 – Theory of Mind Theory	67
Section 2.3.3.1 – Strengths of Theory of Mind Theory	68
Section 2.3.3.2 – Weaknesses of Theory of Mind Theory	68
Section 2.3.4 – Executive Dysfunction Theory	70
Section 2.3.4.1 – Strengths of Executive Dysfunction Theory	70
Section 2.3.4.2 – Weaknesses of Executive Dysfunction Theory	71
Section 2.3.5 – Section Conclusion	71
Section 2.4 - Eliminativism About Autism	71
Section 2.5 – Chapter Conclusion	73
Chapter 3 – A Further Essentialist Theory of Autism: Monotropism	75
Section 3.1 – Chapter Introduction	75
Section 3.2 – Introducing Monotropism	75
Section 3.2.1 – The Extant Account of Monotropism	76
Section 3.2.2 – Lack of Academic Recognition	78
Section 3.2.3 – Terminological Limitations	78
Section 3.2.3.1 – Ambiguities in ‘Interest’	78
Section 3.2.3.2 – Ambiguities in ‘Attention’	79
Section 3.2.4 – Explanatory Limitations	80
Section 3.3 – Specifying the Character of Monotropic Processing	81
Section 3.3.1 – The Monotropic/Polytropic Continuum	81
Section 3.3.2 – The Attention Tunnel	82
Section 3.3.3 – Local Over Global Processing	83
Section 3.3.4 – Unimodal Over Multimodal Processing	84
Section 3.3.5 – Style Preference	85
Section 3.4 – Underexplained Symptoms and Behaviours	87
Section 3.4.1 – Perceptual Differences	88
Section 3.4.1.1 – Hyper- and Hypo- Sensitivities	88
Section 3.4.1.2 – Co-Ordination and Integration Difficulties	90

Section 3.4.1.3 – Enhanced Perceptual Processing	91
Section 3.4.1.4 – Difficulties with Shifting Attention.....	91
Section 3.4.1.5 – ‘Meltdowns’	92
Section 3.4.2 – Cognitive Differences.....	93
Section 3.4.2.1 – Difficulties with Contextual Understanding	93
Section 3.4.2.2 – Difficulties with ‘Mentalising’	94
Section 3.4.3 – Social Differences.....	94
Section 3.4.3.1 – Social Inattentiveness	95
Section 3.4.3.2 – Early Language Regression	96
Section 3.4.3.3 – Keeping Up with Social Cues	97
Section 3.4.3.4 – Difficulties with Social Context	98
Section 3.4.4 – Recommendations for Practice	99
Section 3.5 – Self-reported Experiences	99
Section 3.6 – Methodological Limitations	104
Section 3.6.1 – Narrow Range of Self-Reported Experiences	104
Section 3.6.2 – Lack of Empirical Support	105
Section 3.7 – Chapter Conclusion	106
Chapter 4 – The Problem of Attention	107
Section 4.1 – Chapter Introduction.....	107
Section 4.2 – The Extant Monotropism Account.....	107
Section 4.3 – The Discipline Level Problem	109
Section 4.4 – Chapter Conclusion	112
Chapter 5 – The Selection for Action Theory of Attention	113
Section 5.1 – Chapter Introduction.....	113
Section 5.2 – Wu’s Framework and the Selection for Action Account.....	113
Section 5.2.1 – Introducing Wu’s Framework	113
Section 5.2.2 – Question 1 – The Metaphysical Question	115
Section 5.2.3 - Question 2 – The Function Question.....	121
Section 5.2.4 – Question 3 – The Properties Question	122
Section 5.2.5 – Question 4 – The Mechanism Question.....	123
Section 5.2.6 – Question 5 – The Consciousness Question.....	124
Section 5.2.7 – Framework Overview.....	129
Section 5.2.8 – Assessment of Wu’s Methodology	129

Section 5.3 – Compatibility with Monotropism	130
Section 5.4 – Chapter Conclusion	131
Chapter 6 – The Cognitive Unison Theory of Attention	132
Section 6.1 – Chapter Introduction.....	132
Section 6.2 – The History of the Problem of Attention	132
Section 6.2.1 – William James	132
Section 6.2.2 – The Metaphysical Issue	133
Section 6.3 – The Metaphysical Question	134
Section 6.3.1 – Difficulties with Answering the Metaphysical Question	134
Section 6.3.2 – Addressing the Metaphysical Question.....	134
Section 6.4 – Contending Accounts of Attention	136
Section 6.4.1 – Process-first Accounts of Attention.....	137
Section 6.4.2 – Competing Adverbialist Accounts of Attention	139
Section 6.4.2.1 – Attitude-based Adverbialism.....	139
Section 6.4.2.2 – Disposition-based Adverbialism and Explanation- Based Adverbialism	142
Section 6.4.2.3 – The No-Self-Explaining Limitation	143
Section 6.4.2.4 – Summarising Mole’s Assessment of Competing Adverbialist Accounts	144
Section 6.5 – The Cognitive Unison Theory	144
Section 6.5.1 – Defining ‘Tasks’	145
Section 6.5.2 – Defining ‘Cognitive Processes’	147
Section 6.5.3 – Defining ‘Understanding’	148
Section 6.5.4 – Summarising the Cognitive Unison Account	149
Section 6.6 - Explanatory Power of the Cognitive Unison Theory	149
Section 6.6.1 – Accounting for Task Failure	149
Section 6.6.2 – Mastery and Inattentive Performance	150
Section 6.6.3 – Superfluous Competence	151
Section 6.6.4 – Lack of Privileged Processes	152
Section 6.6.5 – The Privative Character of Cognitive Unison	154
Section 6.6.6 – Divided Attention	156
Section 6.6.6.1 – Easy Cases of Divided Attention	157
Section 6.6.6.2 – Difficult Cases of Divided Attention	158
Section 6.6.6.3 – Quiescence in Divided Attention	159

Section 6.6.6.4 – Divided Attention to Composite Tasks	160
Section 6.7 – Filling in Wu’s Framework	161
Section 6.8 – Compatibility with Monotropism	162
Section 6.9 – Chapter Conclusion	163
Chapter 7 – The Priority Structure Account of Attention	165
Section 7.1 – Chapter Introduction.....	165
Section 7.2 – Arguments Against Reductionism	165
Section 7.3 – The Priority Structure Account of Attention	168
Section 7.3.1 – The Features of Attention	168
Section 7.4 – Attention is an Activity	169
Section 7.5 – Priority Structures	171
Section 7.5.1 – The Structural Elements of Priority Structures	172
Section 7.5.2 – The Structural Relations of Priority Structures	173
Section 7.6 – Attention and Consciousness	174
Section 7.7 – Assessment of the Priority Structure Account	181
Section 7.7.1 – Strengths of the Priority Structure Account	181
Section 7.7.1.1 – Overcoming the Heterogeneity Problem of Attention.....	181
Section 7.7.1.2 – Replacing Capacity Limitation.....	182
Section 7.7.2 – Weaknesses of the Priority Structure Account.....	183
Section 7.8 – Wu’s Criticisms of the Priority Structure Account.....	185
Section 7.9 – Filling in Wu’s Framework	188
Section 7.10 – Chapter Conclusion	189
Chapter 8 – Introducing Priority Monotropism.....	191
Section 8.1 – Chapter Introduction.....	191
Section 8.2 – Introducing Priority Monotropism	191
Section 8.2.1 – Revising Monotropism: Priority Monotropism.....	192
Section 8.2.1.1 – Removing ‘Attentional Resources’ and Capacity Limitation.....	192
Section 8.2.1.2 – Differences in the Function of Attention.....	193
Section 8.2.1.3 – Redefining ‘Interests’	194
Section 8.2.1.4 – Redefining ‘Narrow’ and ‘Broad’	195
Section 8.2.1.5 – Differences in Understanding Inattentional Blindness	197
Section 8.2.1.6 – The Phenomenology of Attention	197

Section 8.2.2 – Summarising the Core Claims of Priority Monotropism	198
Section 8.3 – Resolving the Heterogeneity Problem.....	199
Section 8.3.1 – Addressing the First Prong of the Heterogeneity Problem: Accounting for a Wide Range of Seemingly Disparate Symptoms.....	200
Section 8.3.1.1 – Accounting for Differences in Behaviours and Interests: Restricted and Repetitive Behaviours and Interests	201
Section 8.3.1.2 – Accounting for Communication Differences.....	210
Section 8.3.1.2.1 – Differences in Language Use and Language Development	210
Section 8.3.1.2.1 – Differences in Joint Attention and Joint Action	211
Section 8.3.1.3 – Accounting for Perceptual Differences	212
Section 8.3.1.3.1 – Attentional Exclusion and Inattentional Blindness	212
Section 8.3.1.3.2 – Differences in Multisensory Processing	214
Section 8.3.1.3.3 – Simultaneously Co-Occurring Hyper- and Hypo- Sensory Sensitivities.....	215
Section 8.3.1.3.4 – Differences in the Perception of Global Motion Coherence	215
Section 8.3.1.3.5 – Improved Detail Perception	216
Section 8.3.1.3.6 – Sensory Distress and Meltdowns	217
Section 8.3.1.3.7 – Differences in Interoceptive Awareness	218
Section 8.3.1.4 – Differences in Motor Skills.....	223
Section 8.3.1.5 – Section Conclusion	223
Section 8.3.2 – Addressing the Second Prong of the Heterogeneity Problem: Individual Variance	223
Section 8.4 – Accounting for Aspects That Have Been Identified by Other Essentialist Theories	227
Section 8.4.1 – Theory of Mind	227
Section 8.4.2 – Weak Central Coherence and Enhanced Perceptual Functioning	228
Section 8.4.3 – Executive Dysfunction	229
Section 8.5 – Chapter Conclusion	229
Chapter 9 – Directions for Future Research	231
Section 9.1 – Chapter Introduction.....	231
Section 9.2 – Accounting for Co-Occurring Conditions	231

Section 9.2.1 – Savant Syndrome	232
Section 9.2.2 - Hyperlexia.....	237
Section 9.2.3 - Dyspraxia.....	239
Section 9.2.4 – Agnosia	240
Section 9.2.5 - Alexithymia	241
Section 9.2.6 - Synaesthesia	241
Section 9.2.7 – Mirror Touch Synaesthesia	243
Section 9.2.8 – Stress and Anxiety.....	244
Section 9.2.9 – Somatic Conditions.....	246
Section 9.2.10 – Explanatory Limitations	247
Section 9.3 – Explaining Inconsistencies in Autism Research	248
Section 9.3.1 – Inconsistencies in Local/Global Task Performance	248
Section 9.3.2 – Inconsistencies in Explanations of Expectations and Priors	249
Section 9.4 – Improved Foundations for Future Examination of Empirical Testability and Practical Support	250
Section 9.5 – Chapter Conclusion	250
Conclusion	252
References	254

Acknowledgements

In loving memory of my father, Stewart Annan

I am extremely grateful to you for introducing me to the wonders of the unknown and for encouraging in me the determination and curiosity required to explore them.

I will do my best.

To my primary supervisor, Professor Fiona Macpherson, thank you for believing in me and for helping me to grow. I have improved significantly under your guidance. I could not ask for a better mentor.

To my secondary supervisor, Dr David Simmons, thank you for introducing me to the field of autism research and for helping me to see how interesting and important it is.

I would like to thank the CK Marr Educational Trust for fully funding both my MSc and PhD studies. None of this would have been possible without your generous support.

I would also like to thank the Kennedy Memorial Trust for fully funding my one year visitation to MIT and Professor Alex Byrne for hosting me while I was there. I am very grateful that I was able to undertake this incredible opportunity.

Additionally, I would like to thank the Japan Society for the Promotion of Science for fully funding my summer visit to the University of Tokyo and Professor John O'Dea for hosting my visit.

To my mum, Alison Annan, thank you for your support throughout this journey and for reading my work and for always listening to my ideas.

To my fiancée, Katie Walker, thank you for always believing in me and for being there with me every step of the way. You deserve an honorary degree for the time you've invested in supporting me throughout the development of this thesis.

To Vera and Anne Walker, thank you for supporting me and for reading through my work.

To my best friend, Anders William Knight, thank you for believing in me and supporting me.

I would also like to thank all of my friends, including students and staff at the University of Glasgow, for their support. I wish you all the best.

Author's Declaration

I am aware of and understand the University's policy on plagiarism and I certify that this thesis is my own work, except where indicated by referencing.

Introduction

I have two primary objectives throughout this thesis. The first objective is to develop a novel theory of autism, which I will call ‘Priority Monotropism’. The second objective is to demonstrate the explanatory power of this theory.

In chapter 1, ‘Understanding Autism: Highlighting the Necessity of a New Theory of Autism’, I will demonstrate that the development of a new theory of autism is necessary. I will do this by illustrating that our current understanding of autism is insufficient. In this chapter, I will show how our insufficient understanding of the condition leads to practical harms for autistic people and the wider autism community (including family members and caregivers), as well as how this misunderstanding has hindered autism research efforts. I will explain that, with our current understanding of the condition, it is difficult to implement sufficient diagnostic methodologies and support services.

In chapter 2, ‘The Heterogeneity Problem of Autism: Essentialist Theories and Eliminativism’, I will further motivate the need for a new theory of autism by showing that current theories are insufficient. I will introduce four main theories of autism in this chapter: Weak Central Coherence theory, Theory of Mind theory, Enhanced Perceptual Functioning theory, and Executive Dysfunction theory. I will explain that all of these theories fail to overcome the heterogeneity problem. I will explain that there are two parts to the heterogeneity problem of autism. First, there is an extremely wide range of symptoms associated with autism, between which there is no clear conceptual link. Second, there is a high degree of variance between the symptom profiles of autistic individuals. I will explain why these theories are unable to overcome even the first part of the heterogeneity problem and, thus, why this has led to eliminativism about autism.

In chapter 3, ‘A Further Essentialist Theory of Autism: Monotropism’, I will introduce a further theory of autism, Monotropism. I will demonstrate that Monotropism has a lot of potential explanatory power and that it is better equipped to tackle the heterogeneity problem than other essentialist theories of autism. However, I will explain that Monotropism is insufficient in its current state. This is largely due to the vagueness of the account and the fact that key terms—such as ‘attention’ and the related term ‘interest’—are not clearly explained. This is important as Monotropism is a theory that characterises the core essence of autism through differences in attentional style. I will conclude that while this theory has a lot of

promise, its potential explanatory power has not been sufficiently actualised. Throughout this thesis I will develop a revised account of Monotropism in order to address these problems.

In chapter 4, ‘The Problem of Attention’, I will briefly examine potential reasons why the authors of the extant account of Monotropism did not provide a clear definition for their use of the term ‘attention’. I will explain that this issue is not exclusive to their account.

Difficulties with providing a clear explanation of what attention is reflect a deeper discipline-level problem within psychology. It is because of this issue that, in order to create an improved account of Monotropism, I will first need to consider philosophical theories that more clearly characterise the nature of attention.

I will examine the following three philosophical theories of attention: the Selection for Action theory (chapter 5), the Cognitive Unison theory (chapter 6), and finally the Priority Structure account (chapter 7). It is important to note that these are theories of attention in general and not theories about attention in autism. For each of these theories, I will provide both a general assessment of their core claims and an assessment of their suitability to provide a foundation for a revised account of Monotropism by examining the compatibility between the core claims of these theories and the core claims of Monotropism.

I will conclude that the Priority Structure account of attention provides the best foundation for a revised account of Monotropism. I will demonstrate this in chapter 8, ‘Introducing Priority Monotropism’, by introducing a revised account of Monotropism that uses the Priority Structure account of attention as its foundation, which I will call ‘Priority Monotropism’. I will highlight the key differences between the core claims of Priority Monotropism and extant Monotropism. I will then demonstrate the explanatory power of this theory by showing how it can address both parts of the heterogeneity problem. I will do this first by demonstrating that Priority Monotropism can account for the wide range of symptoms and behaviours in autism (by providing an explanation of these symptoms and behaviours that conceptually links them) and second by showing that it can account for variance in symptom profiles between autistic individuals.

Finally, in chapter 9, ‘Directions for Future Research’, I will further demonstrate the potential explanatory power of Priority Monotropism by briefly explaining how this theory can potentially allow for an improved understanding of other aspects of autism. In this chapter, I will briefly explain how considering autism through the lens of Priority Monotropism might allow for an improved understanding of why autism commonly co-occurs with a range of

other, seemingly unrelated, conditions. I will also briefly consider how Priority Monotropism may illuminate and resolve inconsistencies in autism research findings. Finally, I will claim that Priority Monotropism has made clearer claims about the nature of autism and that, in doing so, it provides an improved foundation for future empirical research and support efforts.

Terminology Disclaimer

Before I begin, it is important to clarify the terminology that I will be using and justify my reasoning for doing so. Within the field of autism research, there is disagreement about which terms are most appropriate to use when discussing autism.

The central disagreement revolves around whether to use ‘person first’ terminology or ‘identity first’ terminology. These terminologies concern the order in which words should be used when discussing autism. ‘Person first’ terminology places the words relating to personhood first; ‘individual’, ‘person’, ‘woman’, ‘boy’, etc., before ‘autism’. Examples of ‘person first terminology’ include ‘individuals with autism’ and ‘girls with autism’, etc. In contrast, ‘identity first’ terminology does the reverse by placing the identity label, ‘autistic’, first. Examples of ‘identity first’ terminology include ‘autistic individuals’ and ‘autistic men’, etc.

Historically, autism researchers have typically utilised ‘person first’ terminology, in an understandable attempt to ensure that autistic people are considered first and foremost to be people. ‘Person first’ terminology has been endorsed by almost half of all professional researchers (Kenny et al., 2015).

However, while I am sympathetic to this reasoning, I will be utilising identity first terminology as this style has been supported by the autism community. Many autistic people feel that person first terminology makes autism seem like a separable addition to themselves (as something that they ‘have’) and also makes it seem more akin to a disease. It is for this reason that many autistic people prefer identity first terminology, which allows us to understand autism as an important and ingrained part of their sense of self, as opposed to an addition (Autistic Self Advocacy Network, n.d.).

Despite this, it is important to recognise that the use of such terminology is a divisive issue, and neither selection will be satisfactory to the entire autism community. For example, many parents of autistic children prefer person first terminology and may view the term ‘autistic’ to be insulting (Autistic Self Advocacy Network, n.d.).

I am aware of this issue, and it is not my intention to insult, but rather to help to represent autistic people and destigmatise the terms that they typically prefer. This is why I will be using identity first terminology throughout this thesis.

I would also like to note that other researchers often use the terms ‘ASC’ (autism spectrum condition) and ‘ASD’ (autism spectrum disorder). While I will not be using these terms, they will sometimes appear in quoted statements throughout this thesis.

Chapter 1 – Understanding Autism: Highlighting the Necessity of a New Theory of Autism

Section 1.1 – Chapter Introduction

In this chapter, I aim to outline the explanatory gap that my thesis will attempt to fill. I will show that, despite the fact that most people have heard of autism and likely have a commonsense notion of the condition, autism remains poorly understood in a foundational and meaningful way. I will discuss the impact that our current understanding of autism has for both autism research efforts and for the autism community. In this chapter, I will show that our current understanding is insufficient and that this leads to insufficient, subjective, diagnostic methods, lack of sufficient support opportunities, and harmful misconceptions about the condition. This motivates the need for a new theory of autism.

Section 1.2 – Showing the Explanatory Gap

Throughout this thesis, I will be developing a new theory of autism. One major issue with this is that it may not be obvious to my readers that such a theory is necessary; after all, most of us are likely to have a commonsense understanding of autism already. You may have autistic friends or family members, or you may know someone who does. You may have seen renditions of autism and autistic people in the media. After all, autism is not an exceptionally rare condition; around 1% of the UK population is diagnosed as autistic (around 700,000 people) (British Medical Association, 2024). Additionally, autism has received a lot of media attention, especially in the UK and the U.S. You may have read news relating to autism, such as reports that autism rates have risen. You may be aware of specific specialist education programmes or therapeutic practices that are aimed at supporting autistic people.

So, the question is: why do we need a theory of autism, if we already have a commonsense understanding of what autism is?

Having this commonsense understanding may make it difficult to engage with the primary purpose of this thesis – the development of a new theory of autism. Why do we need a new theory of autism given that we already have a commonsense understanding and a diagnostic process that has been used to diagnose 700,000 people in the UK? It may seem that there is little need to develop a new theory of autism, as the current model is already functional.

Throughout this chapter, I will argue that our current understanding of autism is insufficient. I will also demonstrate that our insufficient understanding of autism is not a benign issue, as it leads to many serious negative consequences that harm both the autism community and autism research efforts.

The issue is that our broad commonsense understanding of autism is insufficient for furthering our theoretical understanding of the condition and for providing practical support to autistic individuals and their families. Additionally, this commonsense understanding is likely to vary from person to person and allows for a lot of personal interpretation and misinformation, which may be founded on—often harmful—stereotypes. Furthermore, our commonsense understanding is extremely vague and difficult to summarise in any meaningful way. An example of this can be seen from the National Health Service (NHS) website, which serves as one of the main accessible sources of public information concerning health and mental health conditions in the UK. On the NHS webpage for autism, autism is summarised as follows: “Autistic people may act in a different way to other people” (National Health Service, 2025). This summary is extremely vague and unhelpful, and it showcases how difficult it is to summarise autism concisely based on our current understanding of the condition. This is not the only example of a vague general definition of autism. The National Autistic Society (n.d. a) defines autism as follows: “Autism influences how people experience and interact with the world. It is a lifelong neurodivergence and disability.”

Our commonsense understanding of autism seems to boil down to nothing. The main reason for this is that autism is a highly heterogeneous condition, wherein behaviours and symptoms may vary widely from individual to individual. This heterogeneity means that it is very difficult to provide a clear explanation of what autism is that captures all of the common symptoms and behaviours and is applicable across all cases.


Our poor understanding of autism has led to several serious theoretical and practical consequences that have negatively impacted both the autism community and autism research efforts. In the next section, I will consider the theoretical consequences of failing to provide a sufficient account of autism. I will then examine the ways in which this has negatively affected the autism community in section 1.4, where I will consider the practical consequences.

Section 1.3 – Theoretical Consequences

Section 1.3.1 – Autism Research Efforts

Autism has been extensively researched with a dramatic increase in publications within recent years. Fletcher-Watson and Happé (2019) state that

...in 1988, fewer than 200 papers on autism were published, and the total literature amounted to well under 3,000 specialist publications. Now, a search on PubMed for papers with “autism” in the title or abstract returns well over 4,000 results published in 2017 alone – a rate of more than 80 new papers per week. In total, the scientific literature on autism exceeds 60,000 papers as we write. (Fletcher-Watson and Happé, 2019, p.154)



(Figure 1.1 – National Center for Biotechnology Information, 2026)

This number has since risen from 60,000 to over 90,000 (as shown in Figure 1.1) as of January 2026, with over 8,900 published in 2025 alone. Yet despite this exponential growth of autism research efforts, autism remains a poorly understood condition.

Section 1.3.2 – Eliminativism About Autism

The fact that extensive efforts have been made in the field of autism research with little progress on our understanding of the condition has led to an increase in eliminativist ideology about autism.

Autism eliminativism rejects the idea that autism is a real, unified, condition that can be understood in any meaningful way. Autism eliminativists explain our lack of understanding of autism, despite our extensive research efforts, by stating that our research efforts are misguided. In our research efforts, we are aiming to examine and identify a single unified condition, ‘autism’, and since such a condition does not exist (according to autism eliminativists), our research efforts are bound to be unsuccessful at achieving this aim.

As I will show in chapter 2, autism eliminativism is fuelled by the extent of the heterogeneity problem. There is widespread heterogeneity in symptom profiles between autistic individuals. For example, person A and person B may both be autistic but may not have any significant overlap in their symptoms or behaviours. Furthermore, there is also heterogeneity within the range of symptoms and behaviours in general; many symptoms and behaviours are difficult to connect and explain. For example, reduction in/lack of utilisation of eye contact, sensory sensitivities, and difficulties with social understanding.

Another major issue is that the biological aspects of the condition are also poorly understood. There are no clear biomarkers (neural or genetic markers) for autism

... efforts to identify reliable diagnostic biomarkers, meaningful (biological) subgroups, autism-specific genes or neural circuits, and targets for brain-based and psychopharmacological interventions remain disappointingly unproductive. Current candidate biomarkers for autism—such as particular genetic variants, different brain structures, brain functions, and neuropeptides—are not found in all autism cases (poor sensitivity) and they tend to be associated with many other neurodevelopmental disorders and “normal” conditions (poor specificity). (Verhoeff, 2015, p.444, cited in Cushing, 2026, p.209)

Although it is known that autism can be hereditary, no specific genetic marker for autism has been identified. There appear to be no significant markers for the condition at all (Hus and Segal, 2021).

It is important to note that eliminativism has the potential to negatively impact the autism community if it is widely adopted. Eliminativism is potentially harmful in the following three ways:

First, by potentially causing a disruption to the sense of community that surrounds autism – by removing the unified label that connects a wide range of autistic people.

Second, by creating wider social impressions that autism is not a ‘real’ condition, which could also have a practical impact by reducing autism specific support services.

Third, by potentially preventing further research and investigation into autism as a unified condition, which could reveal information that could be used to support autistic individuals.

Chapman (2020) attempts to take a middle-ground approach by acknowledging the issues that are leading to eliminativism while stating that, even if we accept scientific eliminativism, we should not be completely eliminativist about autism, as it is still a useful term for referring to a specific social group. He argues that autism is best understood as a ‘serial collective’, a group in which people are related to each other through their shared experiences in relation to their social environmental contexts.

Chapman (2020) claims that “autism may remain a meaningful and helpful classification even if it lacks scientific validity”(Chapman, 2020, p.799).

However, while Chapman’s approach may help to prevent the first potential harm, by allowing for a sense of community to remain, it does not prevent the other two potential harms. If autism is seen only as a social kind, then this prevents future scientific research. Also, endorsing a version of autism that lacks scientific validity is likely to only further diminish autism’s status as a ‘real’ condition.

I also have concerns about how useful Chapman’s (2020) serial collective notion of autism really is. First, it seems that there is more to autism than a social kind; there is some biological component, as the condition is hereditary. Second, an autism eliminativist could ask: how can we be sure that all of these people in the social collective really do share experiences? For example, if autistic person A is prone to frequent outbursts of extraversion, as a result of their autism, while autistic person B is prone to extended periods of introversion, as a result of their autism, then the way in which they both experience their social contexts is likely to be different. Third, this does not seem to account for instances where autism overlaps with any commonly co-occurring conditions, such as agnosia, dyspraxia, intellectual disability, etc, as these are likely to have an effect on the way that an individual relates to their social context. Co-occurrence with such conditions, or others, could push one out of the category of ‘autism’ on this view. Fourth, this view also seems to have the potential to allow non-autistic people into this category, as long as they relate to the social context in a specific way. For example, someone who has a social anxiety disorder may have a relational experience to their social environment that is similar to that which is typically found in autism; on this view, these individuals would be classified as ‘autistic’.

I will further discuss eliminativism, and how to overcome it, in chapter 2.

Section 1.4 – Practical Consequences

There are also practical consequences of not having a sufficient understanding of autism. I will break these down into three main categories: consequences for support, consequences for diagnosis, and misconceptions about autism.

Section 1.4.1 – Consequences for Autism Support


Our current level of understanding is insufficient for providing adequate support for autistic people. Many autistic people struggle with foundational elements of daily living, such as securing employment and completing educational programmes. This may affect their ability to live independently in society.

Section 1.4.1.1 – Employment

Autistic people are significantly less likely to successfully obtain and maintain employment than non-autistic people. The Buckland Review of Autism Employment provides a dedicated report of employment statistics for autistic individuals.

The most recent of these reports, published in 2024, states that only 3 out of 10 autistic people (of working age) are in employment (Department for Work & Pensions, 2024). This is drastically lower than the average for non-disabled people, who have an employment rate of 8 out of 10, and is even significantly lower than the average employment rate for disabled people, who have an average employment rate of 5 out of 10 (Department for Work & Pensions, 2024).

Figure 1. Autistic people face lower employment rates compared to other disabled or non-disabled people



(Figure 1.2 – Department for Work & Pensions, 2024)

This disparity cannot be explained by differences in IQ rates, as a recent study by Katusic et al. (2021) has shown that nearly half of all autistic individuals have an average or higher IQ, yet only one third are in employment.

Additionally, it cannot be entirely explained by a lack of formal education, as similar unemployment rates are observed in autistic graduates. Only 36% of whom find full-time employment within 15 months of graduating; this number is half that of non-disabled graduates (Department for Work & Pensions, 2024).

The Buckland review states that this is not caused by a lack of interest in working, as many autistic people wish to work. But rather, it is caused largely by a lack of support both within the workplace environment and during the recruitment process (Department for Work & Pensions, 2024). This statement is supported by statistics published in *Autistica* by Gomez and Sheikh (2023), which show that 77% of autistic people reported that they wanted to work and that the majority of autistic people listed workplace support as being an important factor that would allow them to work.

Work-related challenges identified in this report include underprepared employers, outdated attitudes, lack of sufficient autism directed support from employers and career advisors, and

negative interview experiences due to uncertainty about the interview process, which often prioritises one's social skills over their relevant career-related skills (Department for Work & Pensions, 2024).

These factors show how difficult it can be for autistic individuals to secure and maintain employment, as sufficient support is often not provided by employers.

Section 1.4.1.2 – Education

Autistic individuals may face difficulties with education at all stages—from primary school to university.


In 2021, the National Autistic Society published a detailed report of the experiences of autistic children and their families in relation to education. This report found that many autistic children and parents of autistic children were extremely dissatisfied with the level of educational support that they were being offered (as shown in Figure 1.3 below).

This report found that lack of understanding of what autism is within educational settings was a major contributing factor. It reported that despite the fact that 99% of people had heard of autism, autistic people and their family members did not think that the general understanding of autism was sufficient (National Autistic Society, 2021).

More than half of the autistic children who responded for this report stated that having teachers who do not understand them is the worst thing about school. This report states that teachers having an understanding of what autism is will be essential for autistic children to succeed in education (National Autistic Society, 2021).

Additionally, this report states that a better understanding of autism would allow for improvements in peer-to-peer relationships within educational settings (National Autistic Society, 2021).

We heard from parents of autistic children about their personal experiences of education and the support they receive. We found:



In our autistic children and young people's survey:



(Figure 1.3 – National Autistic Society, 2021)

This is demonstrated by the following quotes from participants in this report:

“[support was] limited and not really tailored to autism and its many faces”.

“My child's secondary school experience was awful. The school just tried to make her leave their establishment as this seemed the easiest option for them.”

(National Autistic Society, 2021)

Teachers, even those working within special educational services, have also claimed that they feel that they require a better understanding of autism and improved training in order

to increase their confidence about the support that they can provide to autistic students (Sanz-Cervera, 2017).

A study by Barned et al. (2011) investigated pre-service teachers' understanding of autism and found that many of them held misconceptions about the condition. The results of this study showed that 93% of participants did not know that autism is a developmental condition, 60% of participants wrongly believed that autism is a condition that can be outgrown, and 20% believed that autism could be caused by trauma.

Having a clearer understanding of what autism is would allow us to better understand the needs of autistic people and allow us to develop better support systems, which could improve employment and education rates.

Section 1.4.1.3 – Suicidal Ideation and Suicide Attempts


Additionally, autistic individuals are significantly more likely to experience suicidal ideation and/or attempt suicide than non-autistic individuals. There is a correlation between the age of the individual and the rate of suicide attempts and suicidal ideation. At all life stages, autistic individuals are more likely to experience suicidal ideation or attempt suicide than non-autistic individuals. However, these rates significantly increase in autistic adult populations. The reverse is true of non-autistic populations, where the rate of suicide attempts and suicidal ideation decreases in the adult years.

This is shown in a study by Conner, Ionadi, and Mazefsky (2023), which reported that autistic children and teenagers are more likely to experience both suicidal ideation and to attempt suicide than non-autistic teenagers. This study found that 20% of autistic children and teenagers had experienced suicidal ideation in the past year and that 10% had attempted suicide. This rate was lower at 14.2% for non-autistic children and teenagers experiencing suicidal ideation and 4.5% attempting suicide.

This study found significant differences between autistic and non-autistic adults regarding the rates of suicidal ideation and suicide attempts. In non-autistic populations, the rates of suicidal ideation and suicide attempts dropped significantly in adulthood, while in autistic populations these rates both increased dramatically. Conner, Ionadi, and Mazefsky (2023) reported that 4.8% of non-autistic adults experienced suicidal ideation and that 0.7% attempted suicide, while 42% of autistic adults reported suicidal ideation and 18% attempted

suicide. Additionally, 60% of autistic adults who were diagnosed as autistic in adulthood experienced suicidal ideation.

Conner, Ionadi, and Mazefsky (2023) conclude that autistic adults are 25 times more likely to attempt suicide than non-autistic adults.



(Figure 1.4 - chart created based on findings by Conner, Ionadi, and Mazefsky (2023)).

One explanation for this significant increase in both suicidal ideation and suicide attempts in adulthood is a lack of support for autistic adults. Autistic adults may find it difficult to cope with the increased pressure of living independently in a neurocentric society. They may struggle with issues surrounding gaining and maintaining employment, successfully completing education, being able to support themselves financially, and developing and maintaining social relationships when they do not have access to sufficient support systems (Department for Work & Pensions, 2024).

Section 1.4.1.4 – Victimisation

A meta-analysis by Trundle et al. (2022) reported that autistic individuals were more likely to experience victimisation than non-autistic individuals. They reported on the prevalence of specific types of victimisation in autism populations. They found that the most common form of victimisation was bullying, which was experienced by 44% of autistic individuals, followed by sexual victimisation, which affected 40% of autistic individuals, then child abuse, which affected 14%, and finally cyberbullying, which was experienced by 13%. Additionally, they found that 84% of autistic individuals experienced more than one form of victimisation (Trundle et al., 2022). The authors note that the rates for victimisation through child abuse may be even higher than this, as there was less available data.

Interestingly, Trundle et al. (2022) reported that autistic individuals who have comorbid learning difficulties were less likely to experience victimisation due to the support that they were receiving for these conditions. This demonstrates the importance of the role that support plays in combatting victimisation. Trundle et al. (2022) claim that providing support for autistic people is important for reducing the rates of victimisation that they experience, and they claim that it is important for us to have an improved understanding of autism in order to do so.

Offit (2010, p.29) has claimed that lack of understanding and support has led to several tragic cases of filicide wherein “Some parents, frustrated beyond reason and sanity, have killed their own children”. It is important to note the parental attitudes towards their children and towards autism that are clearly shown in some of these cases. For example, in the murder of twelve-year-old Ulysses Stable, his father stabbed Ulysses to death before calling the police to report the murder by calmly stating “I’ve terminated the life of my autistic child.” (Offit, 2010, p.29). Or in the case of three-year-old Katie McCarron, who was suffocated by her mother, who claimed: “Nothing is going to help and it’s not going to make any difference. [I] just wanted to end my pain and Katie’s pain.” (Offit, 2010, p.30).

Guan et al. (2021) conducted an investigative study into homicide incidents involving autistic children from 2000 until 2019. They reported that the most common perpetrators were parents or guardians and that the most common reason cited was ‘overwhelming stress’. They claim that this demonstrates the importance of providing sufficient support for autistic people and their families.

Section 1.4.2 - Diagnosis

Diagnosis is another area that is significantly negatively impacted by not having a sufficient understanding of what autism is. Thomas et al. (2011, p.202) claim that "... autism is a clinical diagnosis for which there is no reliable diagnostic marker...".

How can we form accurate diagnostic criteria if we do not first have a foundational understanding of the condition for which we are attempting to create diagnostic criteria? How can we decide which diagnostic procedures to utilise if we do not have secure diagnostic criteria? How can we tell if we are correctly diagnosing autism if we do not know what autism is?

These are some of the biggest theoretical concerns surrounding the diagnostic processes of autism. However, these theoretical concerns have very real practical consequences, which can be seen by examining the current diagnostic criteria and processes that are used to diagnose autism today.

Section 1.4.2.1 – Current Diagnostic Criteria

The current diagnostic method for autism is behavioural observation based on symptoms and criteria that are outlined in the DSM-5.

The criteria listed in the DSM-5 are as follows:

...persistent deficits in the following three areas of social communication and interaction: (i) social-emotional reciprocity; (ii) developing, understanding, and maintaining relationships; and (iii) nonverbal communication. In addition, at least two of the following four behaviors should be present: (i) inflexible to changes in routine; (ii) restrictive or fixated interests that may be abnormal in focus or intensity; (iii) hypo- or hyperactivity in response to sensory input or abnormal fixation with sensory aspects of the environment; and (iv) repetitive movements, speech, or use of items. (Lordan, Storni, and De Benedictis, 2021)

Additionally, in order to receive an autism diagnosis, these symptoms should not be better explained by any other comorbid condition, such as intellectual disability (National Autistic Society, n.d. b).

Section 1.4.2.2 – Problems with the DSM-5 Criteria

The DSM-5 identifies autism by identifying a cluster of behaviours. Misdiagnosis may occur as a result of the rigid opening criteria, which state that in order to receive a diagnosis, the individual must have deficits in all three areas: social-emotional reciprocity, developing understanding and maintaining relationships, and non-verbal communication. Due to the heterogeneous nature of autism, it is possible that some autistic people would not fit this description (I will discuss the heterogeneity issue further in chapter 2). For example, autistic individuals who are non-verbal may display improved non-verbal communication over those who are verbal and therefore would not meet the criteria for a diagnosis on the DSM-5.

Additionally, some of these criteria are vague and can be generally criticised. For example, what is meant by difficulty with understanding and maintaining relationships? Does this mean that in order to obtain a diagnosis the individual must have persistent difficulty with making and keeping friendships and romantic relationships? What constitutes a relationship and what constitutes difficulty? Additionally, autistic people have often reported finding it easier to develop and maintain relationships with other autistic people—showing that these difficulties are not always present.

Furthermore, this approach to these criteria has led to a stricter reclassification of autism. Under this stricter classification, individuals who would previously have been classified as autistic may no longer be classified as such. Volkmar and Reichow (2013) reported that, under this stricter classification, individuals with higher IQ scores were less likely to receive an autism diagnosis, despite the fact that differences in IQ do not feature under the diagnostic criteria.

Cushing (2012) argues that the implementation of stricter criteria in the DSM-5 was, at least partly, a financial decision due to the costs incurred by providing support to a growing number of autistic people. These stricter criteria may also function to curb the perception of an ‘autism epidemic’. However, these issues show a major issue with our current understanding of autism—namely that the question of whether or not someone is or is not autistic can boil down to whether or not there are funds available to support them.

Additionally, the addendum which states that these behaviours must not be better explained by any other comorbid condition is highly problematic. Autism commonly co-occurs with a range of other conditions that may affect intellectual development, such as dyslexia, dyscalculia, and dyspraxia. It can be difficult to tell whether or not the symptoms are better

explained by autism or by co-occurring conditions. Additionally, even if these behaviours are better explained by a co-occurring condition, this does not mean that the individual is not also autistic.

Furthermore, the symptoms listed in the DSM-5 can be masked. Meaning that if the autistic individual does not display these symptoms, then they correctly do not warrant a diagnosis under the current DSM-5 criteria with the use of behavioural observation methods. This is because, on this approach, autism is defined by a range of observable behaviours.

Hus and Segal (2021) also cite Carpenter (2013) as stating that the DSM-5 does not successfully account for a wide enough range of symptoms and behaviours of autism. Examples listed include: ‘issues with play/imagination, shyness or social anxiety, language and failure to meet developmental milestones, behavioral problems and/or temper tantrums, and some poor imitation skills.’

Hus and Segal (2021) claim that the DSM-5 criteria of ‘Restricted and Repetitive Interests and Behaviors’ (RRBs) can be difficult to differentiate from other repetitive motions, for example, tics which are common in Tourette's Syndrome. It is for this reason that Tourette's syndrome is frequently misdiagnosed as autism under the current DSM-5 criteria.

I will further discuss the DSM criteria in chapter 2 and co-occurring conditions in chapter 9.

Section 1.4.2.3 – Problems with Behavioural Observation

Behavioural observation is a highly subjective process that can be skewed by a wide range of factors—leading to inconsistent and unreliable diagnoses. This includes misdiagnosis in both directions: false positive and false negative.

Hus and Segal (2021) cite Prizant (2006) as having uncovered that practical issues, such as the setting of the assessment location, caregivers being present during the assessment, and lack of sufficient staff training, could all lead to misdiagnosis when conducting behavioural observation.

There is also the worry that in our reliance on behavioural observation, we are also making the mistake of prioritising the observable features of autism over and above the internal, personal aspects that are central to autistic people and their experiences.

Another serious problem for behavioural observation is observer bias. The observer may be biased, either intentionally or unintentionally, about who is likely to be autistic, and which behaviours should be considered ‘typical’ or ‘atypical’ in relation to making an autism diagnosis. This issue is especially relevant when considering personal factors about the assessed individuals such as gender, ethnicity, culture, age, and whether or not they have any other conditions or disabilities.

Section 1.4.2.3.1 – Gender

Autism is more commonly diagnosed in boys and men than it is in girls and women. A meta-analysis study by Loomes, Hull, and Mandy (2017) of over 50,000 autistic participants reported that the ratio of autism diagnosis is 3:1, three males for every 1 female (although it has been previously reported as even higher than this at a ratio of 4:1). This has led to the idea that autism is a gendered condition, that is more common among boys and men, which has in turn led to gender stereotyping of autism (Russell et al., 2021).

A popular alternative explanation for this gender divide in autism is that autistic girls and women are being underdiagnosed. It is not the case that they are, at least significantly, less likely to be autistic but rather that they are less likely to be detected (Russell et al., 2021). The leading explanation for why autistic girls and women are less likely to receive an autism diagnosis is that they are more likely to utilise masking strategies that hide their observable autism traits and behaviours (National Autistic Society, n.d. c). Masking is a strategy that autistic people can employ where they copy or perform behaviours similar to those of non-autistic people in order to better fit in and draw less attention to themselves. This means that it is more difficult for us to diagnose autistic girls and women since masking autistic behaviours makes autism unobservable, which is obviously problematic for our current diagnostic method of behavioural observation. Autistic behaviours cannot be observed effectively if they are being masked, and therefore an accurate diagnosis cannot be guaranteed.

Underdiagnosis means that many autistic girls and women are not receiving access to support services. Women who do eventually receive an autism diagnosis tend to receive this diagnosis at a later age than boys or men. Hus and Segal (2021) claim that women receive a diagnosis, on average, 9 years later than boys or men and that the mean age for women to receive an autism diagnosis is 29 years old. This is significant as it shows that even those who do

receive a diagnosis are less likely to have access to early years support or support through their school education.

Furthermore, there is evidence that autistic girls and women may need additional support specifically due to their frequent utilisation of masking strategies, which can negatively affect their mental health. Hus and Segal (2021, p.3521) claim that “In continuous attempts to control and mask their symptoms, adolescent girls may show internalizing problems such as anxiety, depressive symptoms, loss of identity and self-esteem, self-injuries, or suicidal thoughts.”

Supporting evidence for the idea that autistic girls and women are underdiagnosed has been demonstrated by the fact that this gender divide lessens when examining autistic populations with intellectual disabilities. Hus and Segal (2021) report a general gendered ratio of 4:1, but note that this ratio is different for autistic people with an IQ of less than 85. When examining populations with an intellectual disability, the ratio drops from 4:1 to 2:1, showing a much smaller gap.

Why is the gender divide so significant, and why do autistic girls tend to have a lower IQ than autistic boys? An explanation is that autistic girls are more likely to be noticed when they have an intellectual disability and less likely to be noticed when they do not. This connects with the general theory that autistic girls are underdiagnosed due to the fact that they are more likely to mask their autistic symptoms and behaviours. Masking behaviours relate to IQ/intellectual ability—autistic girls with an intellectual disability may be less likely to be able to employ masking strategies effectively, leading to higher rates of detection.

Further issues with the subjective nature of our diagnostic criteria can be shown by the fact that girls and women are less likely to check the RRBs criteria. This is due to the fact that girls are more likely to exhibit restricted interests that are deemed to be more socially acceptable. “Girls’ restricted interests tend to centre on animals, soap-operas, celebrities, pop music, fashion, horses, pets, and literature similarly to neurotypical girls, while boys’ focus is on vehicles, trains, traffic lights, etc.” (Hus and Segal, 2021, p.3521). This can lead to these interests not being considered as ‘autistic’ interests and so ignored when making an assessment.

This shows another serious weakness in our current diagnostic criteria—behaviours and interests are judged against an undefined concept of normalcy. What is it for an interest to be

considered normal, restricted, or pervasive? It seems that this is largely up to the individual making the assessment.

The gender divide in autism is much more likely to reflect difficulties with our subjective diagnostic methods than to reflect actual gendered differences in autism prevalence rates themselves.

Section 1.4.2.3.2 – Ethnicity and Culture

It is also important to consider the cultural differences that could lead to missed or misdiagnosis of autism under our current subjective diagnostic measures.

Our categorisation of behaviours that are ‘typical’ and ‘atypical’ is unlikely to uniformly apply across all cultures. For example, in Western cultures we may consider lack of eye contact during conversation to be a sign of autism, which is often socially looked upon as rude or antisocial behaviour. However, in other cultures this may not be the case. Hus and Segal (2021) note that in Chinese culture it is considered rude for children to make direct eye contact with adults. So, in this cultural context aversion of eye contact is not indicative of autism. This is something that could easily be misinterpreted during an autism assessment.

This makes autism difficult to diagnose in certain groups where the expected behaviours may be different. Behaviours may vary across different ethnocultural groups or across different disability cultures. The observer would ideally need to be aware of an extremely wide range of cultural norms and behaviours in order to take them into account when making a diagnosis. Once again, it is difficult to establish a baseline of ‘typical’ behaviour in order to establish what constitutes ‘atypical’ behaviour.

Additionally, observers may be biased, either intentionally or unintentionally, in their observations and/or conclusions. Recent studies have found that white children and those from higher socioeconomic backgrounds are more likely to be diagnosed with autism at an earlier age and with fewer cases of misdiagnosis before correction. This is despite the fact that autism rates have not been found to vary across different racial groups, meaning that misdiagnosis or late diagnosis is more common in non-white groups (Aylward, Gal-Szabo, and Taraman, 2021).

Section 1.4.2.3.3 – Socio-economic Status

Based on epidemiological data, autism may appear to be a condition that is more prevalent among the upper classes, with higher prevalence rates found in populations with higher socio-economic status (Hus and Segal, 2021). However, there is evidence to suggest that autism prevalence rate disparities between socio-economic classes are due to the underdiagnosis of autism in those with lower socio-economic statuses as a result of barriers that they may face along the pathway to diagnosis (Hus and Segal, 2021). This is especially the case for countries without universal health care, such as the United States, where seeking a diagnosis and accessing support thereafter can be expensive.

Based on ASD epidemiological data that came out of northern hemisphere countries, it appeared as a disorder of those in the upper SES layer of society. However, this perception resulted from the disproportionate underdiagnoses of ASD in lower SES communities due to barriers they experienced in access to identification resources. (Hus and Segal, 2021, p.3521)

Other potential barriers include undertrained medical providers caring for those from lower socioeconomic backgrounds, medical providers having more patients and therefore less time with each patient, parents from lower socioeconomic backgrounds not noticing differences as quickly, and potentially lower rates of parental education (Thomas et al., 2011).

Section 1.4.2.3.4 – Age

Age is another significant factor that can be problematic for our current diagnostic methods to account for.

It is significantly more difficult for autistic individuals who do not receive an autism diagnosis in childhood to receive one in their adult years. One reason for this is that many autistic people employ a strategy of ‘masking’ their autism symptoms so that they are better able to blend in with others. This is a process that can be incredibly detrimental to the mental wellbeing of autistic people but which, sadly, many autistic people deem to be necessary in order to attempt to ‘pass’ as non-autistic and navigate social situations (Bradley et al., 2021). Autistic individuals who were not diagnosed in childhood tend to mask more than those who were diagnosed during childhood (National Autistic Society, n.d. d). This is perhaps due to a

lack of access to support (due to not being previously diagnosed) and/or due to attempts to remedy feeling/appearing different to others without a known reason.

This is especially problematic for ‘high functioning’ autistic individuals who are more likely to have been overlooked in their childhood years (Barahona-Corrêa, 2017) yet who may begin to struggle as social demands and independence-related responsibilities increase with age.

Masking hides the observable behaviours that are typically associated with autism and makes it more difficult to diagnose autism through behavioural observation. The problem is that observable behaviours are not the only component of autism; an autistic individual who is masking is still autistic and could still benefit from support, which they are less likely to receive if they are not provided with an autism diagnosis. In fact, autistic individuals who frequently mask their autism may need greater support and may suffer from mental health difficulties, exhaustion, and feelings of isolation as a result of frequent masking (National Autistic Society, n.d. d).

Section 1.4.2.3.5 – Co-occurring Conditions

Autism commonly co-occurs with a range of other conditions and/or disabilities. This can affect the diagnostic process, as it can make it difficult for the observer to determine which symptoms and behaviours relate to autism and which relate to other conditions (Hus and Segal, 2021). This can lead to both underdiagnosis and overdiagnosis of autism.

Hus and Segal (2021) claim that the presence of comorbid conditions can lead to autism symptoms being disguised, making them more difficult for the observer to identify. This can lead to a false negative diagnosis, wherein the autistic person is falsely diagnosed as non-autistic. These individuals may receive a delayed diagnosis later in life. However, some co-occurring conditions make it difficult to receive a diagnosis at all, even a delayed diagnosis, such as hearing impairments, anxiety, and epilepsy (Hus and Segal, 2021).

Additionally, some co-occurring conditions share some of the core symptoms of autism (Hus and Segal, 2021). This greatly increases the potential for false positive diagnoses, wherein non-autistic people are misdiagnosed as autistic. This is also further complicated by our reliance on behavioural observation, as conditions that share typical behaviour profiles will naturally be difficult to distinguish through behavioural observation alone. For example,

atypical eye-contact is also a symptom of ADHD in toddlers, which may lead to a false positive diagnosis of autism (Hus and Segal, 2021). This is especially challenging as gaze direction is one of the main observable behaviours that is characteristic of autism in infants (due to their limited range of expressible social behaviours).

Additionally, Hus and Segal (2021) claim that the presence of a developmental language disorder, characterised by atypical or delayed language development, can lead to missed early diagnoses of autism. This occurs when the symptoms of autism that relate to atypical language development are mistakenly assigned solely to a general developmental language disorder instead. This is especially concerning given that autism commonly co-occurs with atypicalities in language development.

Hus and Segal (2021) also note that dyspraxia is another condition that commonly co-occurs with autism. Dyspraxia is a condition in which an individual exhibits motor skill impairments; this affects their co-ordination and balance. ‘Clumsiness’ and difficulties with balance and motor skills are also considered to be common symptoms of autism. This can make it difficult to distinguish between the two conditions when using behavioural observation as a diagnostic method. This diagnostic issue can lead to:

- False positive diagnosis of autism – where dyspraxia is mistaken for autism
- False negative diagnosis of autism – where autism is mistaken for dyspraxia
- Missed dual diagnosis – where both dyspraxia and autism are present but only one of these conditions is diagnosed, as the symptoms are assumed to belong to only a single condition (either autism or dyspraxia).

Furthermore, autism may co-occur with an intellectual disability, which is marked by an IQ score below 70 points (Hus and Segal, 2021). The problem is that, using our current diagnostic methods, it can be difficult to tell whether or not younger children have an intellectual disability, autism, or both. Hus and Segal (2021) state that this issue is most problematic for assessing children aged 3 and under, as it is difficult to assess their non-verbal cognitive abilities. They claim that it is easier to distinguish the two conditions by ages 4 to 6 when other autism traits may start to become more apparent.

Concerningly, recategorization of intellectual disability has led to a significant increase in the number of intellectually disabled individuals being given a diagnosis of autism (even when

they were not previously classified as autistic under past criteria). The ‘decrease’ in intellectual disability, which was observed as it was recategorised into other conditions, such as autism, accounts for a 64% increase in autism prevalence rates (Hus and Segal, 2021). This perceived dramatic increase in autism prevalence rates has caused panic and given rise to the notion of an ‘autism epidemic’.

I will further discuss co-occurring conditions in chapter 9, where I will provide a brief explanation of how my theory might account for them in future works.

Section 1.5 – Harmful Misconceptions About Autism

Public misconceptions about autism have the potential to cause harm to autistic people and their families. These misconceptions are enabled by our poor understanding of autism, which leaves room for pseudoscience.

In this section, I will discuss the following misconceptions about autism: general misconceptions (section 1.5.1), misconceptions surrounding vaccinations (section 1.5.2), misconceptions about ‘treatments’ for autism (section 1.5.3), supernatural misconceptions (section 1.5.4), and misconceptions based on historic autism theories (section 1.5.5).

Section 1.5.1 – General Misconceptions

The National Health Service (2025) webpage for autism includes a list of misconceptions in attempts to fight harmful misinformation about autism. This includes statements that autism is not a disease, that autistic people can live full lives with friendships, relationships and employment, that autism is not caused by bad parenting, vaccines, dietary issues, and that autism is not infectious.

The fact that there was a perceived need to include this information about misconceptions surrounding autism, when the National Health Service has not provided information about misconceptions for other conditions (such as ADHD or Down’s syndrome), shows that our common-sense understanding of autism is insufficient and may be based on harmful stereotypes (often derived from historic autism research efforts and attitudes).

Additionally, Autistica (2022) commissioned a YouGov poll that surveyed the British public on their opinions about autism. The results of this poll found that nearly 30% of participants

were unsure if autism was curable, 35% of participants still believed autism to be a learning disability, and 39% of participants believed that autistic people lack empathy.

Autism Alliance UK (n.d.) claim that our general “...lack of understanding allows negative impressions of autism to persist – when we don’t understand something, we will overlay our own explanations, which sometimes come from instinct and sometimes from stereotype...”

Section 1.5.2 – Misconceptions Surrounding Vaccinations

In 1998, Andrew Wakefield, who was previously a gastroenterologist in the UK, claimed that the Measles, Mumps, and Rubella (MMR) vaccine was causing children to develop autism (Offit, 2010). He explained his theory of how the vaccine supposedly caused autism at a press conference. He stated that the measles virus found within the vaccine travelled to the intestine, where it caused inflammation. The intestine then became damaged, which allowed harmful proteins to pass through into the bloodstream. This allowed these proteins to travel through the bloodstream to the brain, where they affected the brain, causing autism. However, he never identified the proteins that were supposedly doing this damage (Offit, 2010).

Wakefield’s theory gained widespread coverage in the media. This caused a large-scale public panic concerning the safety of the MMR vaccine, and many parents refused to have their children vaccinated as a result. Offit (2010) reports that just months after Wakefield’s statement, vaccination rates plummeted from 90% to 70% (and in some areas of London to as low as 50%). The result of this was the resurgence of measles, which had previously been eradicated in the UK. In 2002, a small Irish hospital admitted 100 children who were experiencing pneumonia and brain swelling as a result of contracting measles. Three of these children died from these complications (Offit, 2010).

Others began to theorise about treatments that could be used to help children who had received the MMR vaccine to recover from their autism (which had supposedly been caused by it). Some claimed that the MMR vaccine had depleted the children’s stores of vitamin A and that administering vitamin A through cod liver oil was an effective treatment for autism. Offit (2010) cites a paediatrician at the Medical College of Virginia as claiming:

When I give [autistic children] vitamin A in the form of cod liver oil, they get well. In a few days a lot of these children look at me, focus, regain eye contact, and talk about

their box of vision growing. One child's IQ score went up 105 points. (Offit, 2010, p.52)

Large-scale investigations were launched into the MMR vaccine and its potential links with autism. A study at Helsinki University examined the records of two million children, and a similar study at the Boston Medical School examined the records of three million children, in order to assess whether there was a link between autism and the MMR vaccine. Both of these studies found no conclusive evidence for a link between autism and the MMR vaccine (Offit, 2010).

Additional evidence against Wakefield's claim was demonstrated by the fact that Japan had discontinued its use of the MMR vaccine in response to Wakefield's theory. Yet despite discontinuing its use, autism prevalence rates in Japan continued to increase, showing that the MMR vaccine was unlikely to be responsible for the development of autism.

Additionally, the lab that had confirmed the presence of measles in the intestines of the autistic participants in Wakefield's study was brought into question. Those working in the lab were tested on their ability to reliably detect the measles virus. They failed to do so, showing Wakefield's theory was founded on unreliable results (Offit, 2010).

Offit (2010) has shown that there were also important theoretical issues with Wakefield's theory. First, Wakefield claimed that it was the MMR combination that caused autism (through the development of bowel problems). However, studies found that combination (MMR) vaccines had the same rates of bowel problems as separate vaccinations (M)(M)(R). Second, Wakefield had 'shown' in his Lancet paper that only one in five of the children had bowel problems before they 'developed' autism. If his theory were correct, then the causal order should be the other way around. The children should have had bowel problems that led to the development of autism, not autism before bowel problems. Third, if autistic children did have a leaky gut through which harmful proteins could travel into the bloodstream, as Wakefield theorised, then we should also have found proteins from the bloodstream in the gut (the leak should go both ways) – this is not what was found.

Despite this growing evidence against Wakefield's theory, he still retained public support. Many began to accuse those who spoke out against Wakefield of having been paid off by vaccination companies and held firm in their beliefs that autism was caused by the MMR vaccine (Offit, 2010). This ideology was likely empowered by the release of a film about

Wakefield's theory, which was titled 'Hear the Silence'. In this film, Wakefield was portrayed as a public hero fighting back against the 'evil' vaccination manufacturers at all costs.

Deer, an investigative journalist, later reported that Wakefield's research was fraudulent and unethical. He found that Wakefield had lied about where he had recruited his sample of participants from. He had initially reported that the autistic children who had participated in his study had been recruited during unrelated medical examinations that he had carried out. Deer found that these participants had actually been recruited by Barr, a personal injury lawyer, who was representing their parents as part of an anti-vaccination movement (Offit, 2010). Deer also found that Wakefield was paid by Barr to carry out his 'research' in this area without ever acknowledging this in his paper or at conferences (Offit, 2010).

Additionally, Deer found that Wakefield had not been given ethical approval for his study, which involved performing dangerous procedures on his child participants. These procedures involved the use of general anaesthesia, spinal taps, drawing large quantities of blood, taking biopsies, and invasive surgeries. One five-year-old child in Wakefield's study was seriously injured and left in critical condition when his colon was perforated in several places during one of the procedures (Offit, 2010).

Despite this, the idea that vaccines caused autism persisted. On a public broadcast, RFK Jr made the following statement:

[Autism] shatters the entire family. It destroys them. These kids [are] perfectly healthy, wonderful children, and they're brought into these pediatricians' offices, the wonderful family doctor who they trust. And they're given the shot because they're told its going to save your life. And they come out of that pediatrician's office and by late afternoon they are having seizures. Their parents bring them to the hospital and then the kids just disappear. This child who was a complete human being just disappears. He stops talking. He starts banging his head against the wall. He starts biting himself. He stops interacting with his siblings. And these people know this and they are continuing to inject children with this horrible, horrible toxin. (Offit, 2010, p.106)

Offit (2010) claims that the blame shifted from the MMR vaccine to vaccines in general.

Kirby claimed that thimerosal, a mercury-based compound that acts as a preserving agent in vaccines, was responsible for the development of autism due to the high mercury content. In

2002, he successfully fought for the removal of thimerosal from vaccines. Despite the removal of thimerosal from vaccines, six years later autism rates had continued to increase.

Kirby refused to back down in the face of this evidence. He claimed that during this six-year testing period, an increase of mercury must have been introduced into the environment somehow. He blamed coal burning in China, Californian wildfires, and increased rates of cremations in California (which he claimed led to mercury being released from dental fillings) for the increased rates of autism during this time (Offit, 2010). The characterisation of autistic people as having been ‘damaged by mercury’ or as being ‘poisoned’ was socially pervasive during this time and caused a lot of harm to the autism community (Offit, 2010).

Despite overwhelming evidence from investigations that have proven that vaccines do not cause autism, many people (especially parents) continue to believe that they do (Gabis et al., 2022). This belief has empowered an anti-vaccination movement, which has led to the resurgence of dangerous viruses that were previously considered to have been eradicated in the UK (such as the resurgence of measles) (Rao and Andrade, 2011).

Section 1.5.3 – Misconceptions About ‘Treatments’

The fact that autism is not well understood means that it is difficult to develop effective support measures or treatments that can be used to help autistic people who may struggle with certain aspects of their condition.

Parental desperation for a treatment, and in many cases desperation for a ‘cure’, has led to the endorsement of many harmful treatment methods. These methods are often proven to be both ineffective and potentially dangerous, and yet many parents still seek them out.

Section 1.5.3.1 – Harmful Treatments: Chelation Therapy

Chelation therapy arose as a response to the mistaken idea that the mercury found in certain vaccines could be a potential cause of autism. During chelation therapy, autistic individuals would be injected with chemicals that were designed to bind to the mercury that was supposedly causing their autism and remove it (James et al., 2015). Removal of the mercury would supposedly cure the treated individuals of their autism. However, chelation therapy has been found to be both ineffective and harmful (James et al., 2015; Offit, 2010). In 2005, a

five-year-old autistic boy was killed by a doctor who administered a chelation therapy injection to him that stopped his heart (Offit, 2010).

Section 1.5.3.2 – Harmful Treatments: Secretin Injections

Another supposed treatment for autism was secretin injections. Secretin is a hormone that regulates gastric acid and pancreatic bicarbonate (DiGregorio and Sharma, 2023).

The interest in using secretin as a potential treatment for autism began when Horvath et al. (1998) reported improvements in social skills in three autistic children who received secretin during an endoscopy. They claimed that the improvement of social skills in these children was dramatic. As a result of this study, secretin became an extremely popular ‘treatment’ for autism. Offit (2010) claims that within two months, over 2,000 parents had reported trying secretin on their children as a treatment for autism. The demand for secretin vastly exceeded the production rates. This led to predatory practices such as price gouging and fraud. Offit (2010) reports that one doctor paid \$800 for four doses and charged \$8,000 to administer them to patients. Fraudulent liquid versions, containing only water and glycerine, were sold online to many parents.

A study by Sandler et al. (1999) investigated the effectiveness of secretin and found that it was no more effective than a saltwater placebo injection. Parents of children who participated in this study were informed of these results. Despite this evidence, 69% of parents reported that they would still be interested in treating their children with secretin (Offit, 2010).

Section 1.5.3.3 – Harmful Treatments: Facilitated Communication

Facilitated communication claims to provide support for autistic individuals who are non-verbal or who experience communication difficulties. This method pairs the autistic individual with another person who serves as the facilitator, usually someone who is employed to perform this role. The facilitator will hold the autistic person's hand, allowing them to point to letters on a board or a keyboard. This way, they can spell out messages that they wish to convey to others (Braman et al., 1995).

Parents and caregivers were shocked when their children, many of whom had extremely limited communication abilities, were able to express complex ideas and convey profound

emotional experiences through this method. However, parents were unable to achieve the same results when the facilitator was not present; this led to doubts about the legitimacy of this method.

A study by Wheeler et al. (1993) tested the accuracy of facilitated communication by asking participants to identify pictures that were shown to them. Importantly, the pictures were shown separately to the autistic person and the facilitator so that the facilitator would not be able to tell what the autistic person was seeing. If facilitated communication genuinely worked, then the expected outcome would be that the reported word would match the picture that the autistic person had seen, since it was supposedly the autistic person who was leading the communication effort. This study found that out of 180 trials, 0 correct answers were given, providing overwhelming evidence that the facilitators were falsely leading the communication (Wheeler et al., 1993, cited in Braman et al., 1995).

Facilitated communication caused significant harm to the agency of autistic people, putting words in the mouths of those unable to communicate and giving family members false impressions about their personalities and beliefs. Even in the wake of this evidence, the founder of facilitated communication continued to claim that it was a genuine treatment method. He claimed that the findings of this study were due to the fact that studying facilitated communication leads to a breakdown in trust between the facilitator and the autistic individual, which leads to failure of the method (Offit, 2010).

Section 1.5.3.4 – Harmful Treatments: Electroshock Therapy

Autistic individuals may engage in behaviours that are socially unacceptable or that may cause self-harm, such as banging their heads on objects or hitting them with their hands or biting or scratching at themselves (Offit, 2010). These behaviours can be difficult for the autistic person and for their families.

Misunderstandings concerning the voluntariness of these behaviours (and others) in autism have led to treatments focused around providing discipline. One such controversial treatment is electroshock therapy. Dr. Matthew Israel frequently utilised electroshock therapy on the students at the Judge Rotenberg Center in Boston. These therapies were deliberately pain-inducing, as Dr Israel claimed that “If it didn’t hurt it wouldn’t be effective [...] It has to hurt enough so that the student wants to avoid showing that behavior again.” (Offit, 2010, p.29). The use of electroshock treatment stems from the misunderstanding that these ‘problematic’

behaviours are performed voluntarily and that autistic people can be conditioned into not performing them.

Section 1.5.4 – Supernatural Ideology About Autism

Harmful supernatural ideology about autism has been prevalent throughout history and can even still be observed in modern times. These supernatural ideologies about autism function to bridge the explanatory gap that exists in our understanding of the condition. Throughout history, and even still to a lesser extent, autism is viewed as a mysterious condition, which has led to fearmongering about its nature. These supernatural explanations often characterise autism as demonic possession or as being caused by the devil (Ouhtit et al., 2015).

This is obviously harmful to autistic people and their families, as this causes reactions of fear and contempt from those who hold such beliefs.

Kanner (1964) quotes Martin Luther's reported account of a child whose behaviour suggests that he might have been severely autistic [...]. Luther recommended that the child be taken to the nearby river and drowned, because he was possessed of the Devil and had no soul. (Wing, 1997, p.14)

Section 1.5.4.1 – Supernatural Ideology: Changelings

There are tales in folklore that are thought to describe autism and autistic regression, a period of time during infancy where there is perceived regression and autism symptoms become noticeable.

One such story from folklore is the tale of the 'changelings'. This story was popular in certain European countries (Scotland, Ireland, England, and Scandinavia) from the Middle Ages until well into the Victorian Era. According to the story, the devil (or sometimes trolls or fairies depending on the regional variant) would steal infants and replace them with their own malicious child. These replacements would exhibit strange and distressing behaviours, such as screaming, not making attempts to communicate, and not responding to contact from their mothers. The mothers would notice these behaviours and 'know' that these were not their children, and in response they would grieve their losses and try to find their lost children (Waltz, 2013).

Leask and Silove (2005) have claimed that this tale served a practical role—it provided an explanation for why the child seemed different during this period of perceived regression, which is common in autistic children. However, there were serious consequences from this folk tale, as it was firmly believed to be true by many. Many children were abandoned or murdered as a result (Waltz, 2013). If autism had been better understood at this time, then this explanation for autistic behaviours would likely never have emerged.

Section 1.5.4.2 – Supernatural Ideology: Demons and Possession

Superstitions about autism can also be found in modern times and still persist today. In 2003, the belief that an eight-year-old boy's autism was caused by demonic possession led to him being subjected to an exorcism at Faith Temple Church in Wisconsin. During this exorcism, the preacher pressed his knee into the boy's chest for two hours, killing him. The preacher later stated that "We were just praying for him and asking God to deliver him from the spirit that he had" (Offit, 2010, p.30).

Additionally, supernatural beliefs about autism still persist across cultures. For example, in certain regions of Africa, there is a common cultural belief that autism is caused by a curse, witchcraft, or that it represents a dark omen. It is also believed to be transferable to others. As a result, there is significant social stigma around the condition, leading to social isolation of autistic people and their families. This has also led to the underdiagnosis of autism in these areas, as parents may refuse to have their child diagnosed in order to protect them, and their wider family, against social ostracism (Kantawala et al., 2023).

Having an unsatisfactory explanatory account of autism leaves room for such supernatural ideology to continue to exist in modern times. This can further increase the social stigma surrounding autism and cause social and physical harm to autistic people and their families.

Section 1.5.5 – Misconceptions Based on Historic Autism Theories

Many misconceptions about autism have been spurred on by early theories of autism that have influenced public perceptions of the condition and caused pervasive harm. These theories were often formed based on little to no evidence and without considering the perspectives of autistic people or their families.

Section 1.5.5.1 – Harmful Theories: Refrigerator Mother Theory

In 1967, Bettelheim offered the first explanation of autism that offered a ‘cure’ – the Refrigerator Mother theory of autism (Offit, 2010). According to this theory, autism developed as the result of bad parenting practices, specifically as the result of a ‘cold’ parenting style that was adopted by the infant’s ‘refrigerator’ mother (Severson, Aune, and Jodlowski, 2008).

Infants, if totally deserted by humans before they have developed enough to shift for themselves, will die. And if their physical care is enough for survival but they are deserted emotionally, or are pushed beyond their capacity to cope, they will become autistic. (Bettelheim, 1967, p.348, cited in Waltz, 2013, p.73)

This theory offered a practical solution that could be used to ‘cure’ autism, which was to remove autistic children from their mothers. This way they would be able to be placed in more suitable ‘warmer’ care environments where they could ‘thaw’ out and recover from being autistic (Offit, 2010).

This theory obviously caused significant harm to mothers of autistic children who were blamed for their ‘bad’ parenting. This led mothers to experience a lot of shame about their child's autism (Offit, 2010). It also caused harm to the autistic infants and their families when they were removed from their homes.

Bettelheim claimed that he had successfully cured over forty autistic children. The fact that Bettelheim had seemingly been able to treat autism made this theory popular, as no one else had ever been able to achieve results like this (Severson, Aune, and Jodlowski, 2008). The Refrigerator Mother theory gained widespread public exposure in the media; it featured in many major newspapers, including the Times and the New Yorker (Waltz, 2013).

Bettelheim’s work was accessible to the general public and served as a first introduction to autism for many members of the public. This meant that much of the public's understanding of autism around this time was based on the Refrigerator Mother theory.

However, it was later discovered that Bettelheim had falsified these results (Offit, 2010) and that he had even allegedly abused the autistic children who had been placed in his care (Severson, Aune, and Jodlowski, 2008).

Section 1.5.5.2 – Harmful Theories: Extreme Male Brain Theory

Baron-Cohen (2002) characterised autism as the ‘extreme male brain’. This theory was developed in response to the findings that autism is more frequently diagnosed in boys and men than it is in girls and women. However, a modern alternative explanation for this is that girls and women are underdiagnosed due to frequent utilisation of masking strategies and observer biases (see section 1.4.2.3.1).

According to Baron-Cohen, the male brain prioritises ‘systemising’ while the female brain prioritises ‘empathising’. The ‘extreme male brain’ is characterised by hyper-developed systemising and hypo-developed empathising (Baron-Cohen, 2002). Systemising is the process of “developing rules or principles to facilitate the understanding of complex systems (e.g., machinery, weather patterns) [which] underwrites capacities to not only understand but also manipulate objects in one’s physical environment.” (Baron-Cohen, 2003, p.3, cited in Krahn and Fenton, 2012, p.94). While empathising is the process of “perceiving, and emotionally responding to, various mental states in others [which] underwrites capacities to understand and interact with individuals in one’s social environment” (Baron-Cohen, 2003, p.2, cited in Krahn and Fenton, 2012, p.94).

The Extreme Male Brain theory is controversial for a few reasons. First, it has reinforced opinions that autism is a male condition. This has been a serious roadblock to supporting girls and women in receiving an autism diagnosis, enabling them to be further overlooked. Second, it suggests that men and women have distinct cognitive profiles. Krahn and Fenton (2012) argue that Baron-Cohen’s attempt at creating gendered cognitive profiles is unconvincing and that there is no reason for differences in ‘systemising’ and ‘empathising’ abilities to have been gendered in this way.

These gendered differences can also have practical and political implications, as Baron-Cohen uses this theory to argue that it is because of biological differences that fewer women are working in STEM-related areas (Sample, 2012).

The Extreme Male Brain theory has also fuelled harmful stereotypes about autism and led to the development of harmful treatments. Since, according to this theory, autism is characterised by extreme male qualities, it has been hypothesised that autism could be treated by making autistic individuals ‘less boy-like’. Geier proposed that administering Lupron to autistic individuals would shut down the synthesis of testosterone, thus curing them of their autism (Offit, 2010). This ‘treatment’ was also related to chelation therapy. Greier argued

that, since testosterone binds to mercury, by removing the testosterone that the mercury would also be removed (Offit, 2010).

Offit (2010, p.99) claims that Geier made the following public statement on Radio Liberty in 2006: “If you look at [autistic] children, they have high testosterone, they masturbate at age six, they have mustaches, they’re aggressive, and you can treat them by lowering their testosterone and removing the mercury.”

Section 1.5.5.3 – Mistrust of Autism Research Efforts

Unsurprisingly, as a result of these past theories and supposed ‘treatments’, many autistic people are unwilling to participate in autism research efforts. These historic theories and treatments have caused significant harm to the autism community, and as a result, many autistic individuals and their families do not trust that autism researchers have their best interests at heart. Additionally, the majority of historic studies on autism have not considered the perspectives of or consequences for autistic people in the development of their accounts.

This has sparked the ‘nothing about us without us’ movement, which calls for all studies on autism to include autistic researchers, consult autistic individuals, and to utilise qualitative self-reported information where possible.

Section 1.6 – Chapter Conclusion

In conclusion, throughout this chapter I have shown that despite the fact that most people will have heard of autism, it remains a poorly understood condition. I have shown that current support efforts are insufficient and that many autistic individuals are struggling to secure employment or complete educational programmes and are often misunderstood by others. I have provided examples that show that our poor understanding of autism has fostered harmful misconceptions about autism and autistic people, many of which can still be seen today. In doing so, I have shown the importance of developing an improved understanding of autism. In the next chapter, I will explain why autism is difficult to understand by demonstrating the extent of the heterogeneity problem. I will discuss four prominent essentialist theories of autism and explain that none of them are able to successfully overcome the heterogeneity problem.

Chapter 2 – The Heterogeneity Problem of Autism: Essentialist Theories and Eliminativism

Section 2.1 – Chapter Introduction

In the last chapter, I have shown that there is no clear consensus on how we should understand autism. This has led to many different interpretations of the condition throughout its history, many of which have been significantly harmful to autistic individuals and their families. In this chapter, I will explain why it has been, and still is, so difficult to provide an answer to the question: ‘what is autism?’. Why have attempts to answer this question been incredibly vague?

The main reason for this is that autism is a highly heterogeneous condition. I will refer to this as the ‘heterogeneity problem’. The heterogeneity problem of autism has two prongs:

First, autism is a condition with a wide range of seemingly disconnected symptoms. It can be difficult to provide a sufficient explanation for how these symptoms conceptually fit together.

Second, there is often individual variance between autistic individuals. Autistic individuals often have varied symptom profiles, sometimes with very little to no overlap. It is therefore difficult to provide an explanation of what they have in common.

One way to tackle the heterogeneity problem would be to provide an essentialist account of autism – a theory that identifies the core ‘essence’ of autism and uses this essence to answer the question ‘what is autism?’. In order for such an account to be successful, it would need to identify an essence that can explain the wide range of symptoms in autism, and that can also account for individual variation between autistic individuals. This is no small task. Despite the difficulty of this challenge, several essentialist theories of autism have been put forward. In this chapter, I will introduce these theories and show that they are unable to overcome the heterogeneity problem. I will then explain that their failure to do so is motivating eliminativism about autism.

I will conclude by claiming that the failure of these essentialist theories and the rising support for eliminativism motivate the need for a new theory of autism—one that can more effectively tackle the heterogeneity problem.

In the next chapter, I will introduce a further essentialist theory, Monotropism, which has more potential to address the heterogeneity problem. However, I will show that it is

unsuccessful in its current form. After that, and throughout the rest of this thesis, I will develop a novel account, Priority Monotropism, which is better able to tackle the heterogeneity problem.

Section 2.2 – The Heterogeneity Problem

As I have already discussed in the introduction to this chapter, there are two parts to the heterogeneity problem.

The first part of the problem is that there is an extremely wide range of symptoms and behaviours that are common to autism. These symptoms and behaviours are very diverse in their nature and are difficult to conceptually link. Why is it the case that these varied and seemingly unrelated symptoms occur together? These symptoms and behaviours cover a wide range of aspects; there are symptoms that are related to social aspects, cognitive aspects, perceptual aspects, motivational aspects, and behavioural aspects. In section 2.2.1, I will provide a list of symptoms and behaviours that are common to autism, in order to demonstrate this diversity.

The second part of the problem is that there is individual variance between autistic individuals. This is to say that autistic individuals do not always share the same symptom profiles. There can be varying degrees of profile difference between individuals; some individuals may share many symptoms, some may share very few symptoms, and some may not seem to share any symptoms at all. There is a popular saying within the autism community that reflects this aspect of autism: ‘when you’ve met one person with autism, you’ve met one person with autism’ (Cushing, 2026, p.212). This makes it difficult to explain what autistic individuals have in common and why we should identify them all as having a shared condition.

One way to tackle the heterogeneity problem is to provide an essentialist account of autism that identifies a specific core ‘essence’ that is capable of resolving the heterogeneity problem. In order to do so, it must address both parts of the heterogeneity problem. The essence that is identified must be able to provide a way of conceptually linking the wide range of symptoms and behaviours that are commonly found in autism. It must also be able to explain what it is that autistic individuals have in common, even in cases where they do not share many symptoms and behaviours. Additionally, an ideal essentialist account should also be able to provide an explanation for why this individual variance occurs in the first place.

In the next subsection, I will discuss the wide range of symptoms and behaviours that are common to autism before going on to introduce essentialist theories of autism in section 2.3.

Section 2.2.1 – Listing the Wide and Varied Range of Symptoms and Behaviours in Autism

The National Autistic Society (n.d. e) provides an extensive list of common symptoms and behaviours that are found in autism. They divide this list into three main categories: communication differences, differences in behaviours and interests, and sensory processing differences. They also provide an additional list of symptoms and behaviours that are related to stress and difficulties with fitting in.

In this section, I will discuss this list of symptoms and behaviours whilst adding further common symptoms and behaviours from additional sources.

Section 2.2.1.1 – Symptoms and Behaviours Related to Differences in Communication

The National Autistic Society (n.d. e) lists many symptoms and behaviours that are related to differences in communication. ‘Communication’ is a broad category that encompasses aspects of social behaviour, language development, and social understanding, and as such there are many symptoms and behaviours that fall into this category.

The list of symptoms and behaviours that relate to communication differences is as follows:

- limited, delayed, or no speech (including loss of language sometime after learning to speak and avoidance of talking, despite being able to)
- using repeated words and phrases (echolalia), made-up words, technical or very literal speech, pauses, and sophisticated or advanced language
- speaking in an unusual accent
- differences in rate of speech (fast/slow) and intonation (monotone/varied)
- using fewer gestures or facial expressions to communicate
- preferring not to make eye-contact or finding this uncomfortable
- finding it difficult to recognise or read body language or facial expressions

- finding it hard to judge pauses, turn-taking, or to be interrupted during conversations
- being confused by metaphors and idioms, not always understanding hidden meanings or inference and taking phrases literally
- not engaging in a two-way conversation (for example, not asking other people questions about themselves)
- mismatching non-verbal and verbal communication (for example, saying they are feeling an emotion, but with body language that doesn't appear to align with that emotion)
- seems less aware of others around them – for example, they might not respond to their name being called
- autistic babies and toddlers may not smile when smiled at by others or point when they want to show others something or find something interesting
- find following 'social rules' confusing
- talking in depth about their own interests, but not engaging in others' interests
- being direct and honest in communicating opinions (often more direct and honest than people expect; for example, saying they don't like something when they might have been expected to tell a 'polite' lie)
- difficulty telling the difference between someone being friendly or joking and someone trying to bully or hurt their feelings
- difficulty imagining experiences they have never had
- having a clear idea of right and wrong, and a strong sense of justice, but may struggle with grey areas or unclear rules.

Additional communication-related symptoms, not contained within this list, are:

- social interactions and responses being perceived by non-autistic individuals as lacking in empathy (Bellon-Harn et al., 2026; Alkhaldi et al., 2021)
- idiosyncratic use of language (Volden and Lord, 1991)

- and reduced utilisation of protodeclarative pointing (pointing in order to share experiences) with intact utilisation of protoimperative pointing (pointing in order to request something). (Baron-Cohen, 1989)

Section 2.2.1.2 – Symptoms and Behaviours Related to Differences in Behaviours and Interests

The second category that is outlined by the National Autistic Society (n.d. e) is symptoms and behaviours that are related to differences in behaviours and interests.

The list of symptoms and behaviours that relate to differences in behaviours and interests is as follows:

- a strong preference for routine
- distress when a routine changes
- enjoying order and benefiting from clearly defined rules
- making repetitive movements when excited or upset, for example flapping hands, or making the same noise repeatedly
- using repetitive behaviours as a way of staying calm
- having strong passions and dedicated interests (sometimes known as ‘special interests’) for specific subjects/hobbies, that you/they can engage with intensely
- find engaging in their passions and dedicated interests calming and soothing
- finding it easy to focus on specific interests but also find it difficult to disengage and move on to other tasks.

Additional symptoms and behaviours related to differences in behaviours and interests, which are listed by the Centers for Disease Control and Prevention (CDC) (2024), are as follows:

- Lines up toys or other objects and gets upset when order is changed
- Repeats words or phrases over and over (called echolalia)
- Plays with toys the same way every time
- Is focused on parts of objects (for example, wheels)

- Gets upset by minor changes
- Has obsessive interests
- Must follow certain routines
- Flaps hands, rocks body, or spins self in circles

Section 2.2.1.3 – Symptoms and Behaviours Related to Sensory Differences

The third category that is outlined by the National Autistic Society (n.d. e) is symptoms and behaviours that are related to sensory differences.

The list of symptoms and behaviours that relate to sensory differences is as follows:

- being much more or much less sensitive to visual stimuli, sounds, smells or touch than other people
- seeking out sensory inputs, including visuals, sounds, smells or tactile stimuli and using them as a comfort
- showing strong distress, aversion or actively avoiding sources of sensory difficulty
- finding certain sounds, smells, feelings or foods particularly calming or enjoyable, such as coloured lights or being tucked tightly into bed.

Additional sensory-related symptoms, not contained within this list, are:

- Co-occurring hyper- and hypo- sensory sensitivities (Bogdashina, 2016)
- Differences in multisensory processing (Bogdashina, 2016)
- Finding sensory inputs overwhelming in certain situation but not in other situations (National Autistic Society, n.d. f)
- Differences in interoceptive awareness (National Autistic Society, n.d. f)

Section 2.2.1.4 – Symptoms and Behaviours Related to Stress

The fourth category that is outlined by the National Autistic Society (n.d. e) is symptoms and behaviours that are related to stress. It is possible that many of these symptoms and behaviours could be explained by the prevalence of other symptoms and behaviours, such as

sensory differences or communication differences, which may create additional challenges for autistic individuals.

The list of symptoms and behaviours which relate to stress is as follows:

- distressed behaviour that may present as aggression (towards yourself/themselves or others)
- overwhelm and/or meltdowns. ‘Meltdowns’ are characterised by loss of behavioural control in response to overwhelm, which often involves ‘explosive’ behaviours such as screaming, crying, physically lashing out, or engaging in self-injurious behaviours such as hitting oneself (National Autistic Society, n.d. g)
- masking in social situations and needing to release the stress of doing so after returning to a ‘safe space’. ‘Masking’ is the process of effortfully attempting to suppress autism symptoms in order to blend in with others and present oneself in a way which is more ‘socially acceptable’ (National Autistic Society, n.d. d)
- difficulties in school/workplace
- appearing withdrawn or finding it difficult to engage with others when recovering from social situations (sometimes known as ‘shutdown’)
- being misunderstood by others
- reaching a point of breakdown, severe mental health difficulties or burnout.
- feeling overwhelmed by anxiety or anger and having difficulty keeping calm in stressful situations
- experiencing distress or panic
- having difficulty transitioning between classes, schools and further education
- having difficulty engaging with school or being unable to attend school full time
- having difficulty remaining in employment due to problems maintaining relationships with colleagues and managers or because some workplaces or offices can be overwhelming environments
- struggling with self-harm or suicidal ideation
- feeling lonely and isolated
- struggling with everyday tasks, such as cooking, cleaning, self-care, finances and life admin.

Section 2.2.1.5 – Symptoms And Behaviours Related to Difficulties with Fitting In

The final category that is outlined by the National Autistic Society (n.d. e) is symptoms and behaviours that are related to difficulties with fitting in. Again, many of these symptoms and behaviours could be explained by the prevalence of other symptoms and behaviours, such as differences in communication.

The list of symptoms and behaviours that relate to difficulties with fitting in is as follows:

- struggling with things that peers don't (for example, socialising or daily tasks)
- experiencing things differently from peers
- feeling like an 'alien' and/or they don't fit in with other people
- having to put effort into behaviours that others seem to do automatically (for example, forcing yourself/themselves to make eye contact)
- completing tasks in unusual or creative ways
- using coping strategies to make it through the day
- being excluded, bullied, or not fitting in with others
- finding it difficult to have meaningful relationships with people
- finding it more difficult than most to answer interview questions, write CVs/applications and find employment.

Section 2.2.1.6 – Symptoms And Behaviours Related to Differences in Motor Skills

The National Autistic Society (n.d. h) also lists symptoms and behaviours that are related to both fine and gross motor skills:

Differences in gross motor skills may affect the individual's ability to walk, run, ride a bike, and to throw and catch objects.

While differences in fine motor skills involve differences in balance, eye motor control (this is the ability for our eyes to work together in order to follow and/or hold an object in the line of vision as needed), hand–eye coordination, manual dexterity, hand and finger strength, body awareness (proprioception), bilateral coordination –

using two hands together with one hand leading (e.g. cutting with scissors while holding the piece of paper).

These differences in fine motor skills may affect the individual's ability to use tools effectively, write, dress themselves etc. (National Autistic Society, n.d. h)

Section 2.2.2 – The Heterogeneity Problem and the Difficulty of Defining Autism

In this section I have shown that autism has an extremely wide range of seemingly unconnected symptoms and behaviours that are difficult to conceptually link. It is difficult to explain, for example, what not making eye contact, sensory seeking, delays in language development, and preference for routine all have in common.

Additionally, as previously discussed in section 2.2, not all autistic individuals share the same profile of symptoms. For example, one individual may frequently utilise stimming behaviours such as hand flapping but may have a typical utilisation of language, while another individual may not utilise stimming behaviours but may have an atypical utilisation of language.

This problem has implications for the diagnostic criteria for autism. This makes it difficult to determine which symptoms and behaviours should be considered to be indicative of autism.

Section 2.2.2.1 – The DSM-5 and the Cluster of Behaviours Model

The DSM-5 uses the 'Cluster of Behaviours' (CoB) model. In order for a diagnosis of autism to be warranted, the individual must have a certain number of symptoms from across a certain number of categories. These are listed as follows:

The criteria in the DSM-5 for diagnosing ASD include 3 listed deficits in social communication and social interactions. Clinicians must be sure that these characteristics are not due to developmental delay alone. To be diagnosed with ASD, an individual must meet all three of the following criteria:

Difficulties in social emotional reciprocity, including trouble with social approach, back and forth conversation, sharing interests with others, and expressing/understanding emotions.

Difficulties in nonverbal communication used for social interaction including abnormal eye-contact and body language and difficulty with understanding the use of nonverbal communication like facial expressions or gestures for communication.

Deficits in developing and maintaining relationships with other people (other than with caregivers), including lack of interest in others, difficulties responding to different social contexts, and difficulties in sharing imaginative play with others.

The criteria in the DSM-5 also include demonstrating at least 2 of the following 4 restricted and repetitive behavior, interests, or activities:

Stereotyped speech, repetitive motor movements, echolalia (repeating words or phrases, sometimes from television shows or from other people), and repetitive use of objects or abnormal phrases.

Rigid adherence to routines, ritualized patterns of verbal or nonverbal behaviors, and extreme resistance to change (such as insistence on taking the same route to school, eating the same food because of color or texture, repeating the same questions); the individual may become greatly distressed at small changes in these routines.

Highly restricted interests with abnormal intensity or focus, such as a strong attachment to unusual objects or obsessions with certain interests, such as train schedules.

Increased or decreased reactivity to sensory input or unusual interest in sensory aspects of the environment, such as not reacting to pain, strong dislike to specific sounds, excessive touching or smelling objects, or fascination with spinning objects. (Children's Hospital of Philadelphia Research Institute, 2020)

This approach does not require the identification of an essence, but rather it is simply a practical, non-explanatory measure.

This approach faces several problems: first, it is unclear how these symptoms and behaviours are linked (if at all). This leads to questioning about whether these symptoms and behaviours really should be indicative of autism, as without an explanation of how they are related, this list seems arbitrary. This problem is further fuelled by the fact that this range of criteria is constantly being updated and revised with every new edition of the DSM, and that some of these updates have been motivated by practical social and financial considerations (as opposed to being solely motivated by attempting to identify autism).

As I have previously discussed in chapter 1, Cushing (2012) claims that changes to the diagnostic criteria can restrict who receives a diagnosis of autism and that restrictions may even be intentionally applied for practical reasons. Cushing (2012, p.32) quotes Dr David J. Kupfer, who was chairman of the task force responsible for making the revisions for the fifth edition of the DSM, as follows: "We have to make sure not everybody who is a little odd gets

a diagnosis of autism or Asperger disorder [...] It involves a use of treatment resources. It becomes a cost issue.”

Second, this approach allows for individual variation in symptom profiles but does not provide an explanatory account for such variation. Using the CoB approach allows for person 1 to warrant a diagnosis of autism with potentially very little to no overlap between their symptoms and the symptoms of person 2 (who also warrants a diagnosis of autism based on the criteria outlined). However, Cushing (2012) criticises the CoB approach for being unable to explain why we should consider both of these individuals to be autistic or to share a condition.

The CoB approach is to identify key *types* and demand that each be instantiated by one or more *tokens*. No single individual characteristic is essential, but each type must be represented. However, what this list gains in flexibility, it loses in specificity. Now it seems that a wide variety of differing individuals could all meet the same diagnosis. [...] Why should we assume that there is some essentially similar syndrome that they share? (Cushing, 2012, p.31)

In the next section, I will introduce several popular essentialist theories of autism and assess their ability to overcome the heterogeneity problem.

Section 2.3 – Essentialist Theories of Autism

In this section, I will examine several essentialist theories of autism. I will discuss the core ‘essence’ that is identified by each theory and show the strengths and limitations of each account. However, I will ultimately argue that none of these theories can successfully overcome the heterogeneity problem.

The theories that I will discuss in this section are: Weak Central Coherence (WCC) theory, Enhanced Perceptual Functioning (EPF) theory, Theory of Mind (ToM) theory, and Executive Dysfunction theory.

Section 2.3.1 – Weak Central Coherence Theory

Weak Central Coherence (WCC) is a theory of autism that characterises the essence of autism as a local processing bias (caused by an impairment in global processing abilities) that may lead to difficulties with integrating information and with gestalt perception (Happé and Frith, 2006).

The terms ‘local’ and ‘global’ can be briefly defined as follows:


Local processing – narrowly distributing perceptual attention in order to process stimuli in a detail-orientated way.

Global processing – widely distributing perceptual attention in order to process stimuli as wholes, attending to the overall Gestalt as opposed to the constituent (local) parts (Ji et al., 2019).

This definition shows that the term ‘local-level stimuli’ is used to describe stimuli that are smaller parts of larger respective wholes. For example, a lampshade may be a local-level stimulus that is a smaller part of a lamp (that is itself a respective global whole). Global-level stimuli are, respectively, larger stimuli that consist of constituent local-level parts. For example, a lamp may be considered a global-level stimulus as it consists of many constituent local-level stimuli, such as a lampshade, lightbulb, wire, etc.


The central claim of WCC is that autistic individuals have a reduced ability to cohere local-level stimuli into global wholes. This means that their perception of the environment consists of fragmented local-level stimuli (de Jager and Condry, 2020). This fragmented perception makes it difficult for them to ‘see the big picture’ or get the general ‘gist’ of a situation or scenario.

WCC is typically evidenced by improved performance of autistic participants on local-over-global processing tasks, such as embedded figures tasks, Navon tasks, and visual search tasks. Embedded figures tasks require the participant to identify smaller shapes embedded within larger shapes (focusing on the ‘local’ aspects of the task and disregarding the ‘global’ aspects). An example of an image from an embedded figures task can be seen below in Figure 2.1.




(Figure 2.1 – Happé, 2013 a)

Navon tasks require the participant to identify small characters (typically letters or numbers) that are contained within larger characters (again, focusing on the local-level stimuli and disregarding the global-level stimuli). An example of a Navon task image can be seen in Figure 2.2 below.



(Figure 2.2 – Guy et al., 2019)

Visual search tasks require the participant to identify specific targets amongst a range of distractors (focusing on the local level target and disregarding the global level pool of distractors). Two examples of images used in a visual search task can be seen below in Figure 2.3.



(Figure 2.3 - Erez, Shahar, and Katz, 2006)

Section 2.3.1.1 – Strengths of Weak Central Coherence Theory

WCC theorists claim that this local-over-global processing bias can account for a range of symptoms in autism, including insistence on sameness, attention to parts of objects, and uneven cognitive profile (Booth & Happé, 2010). Additionally, Cushing (2012) claims that WCC can successfully provide an explanation of several symptoms and behaviours that are common to autism. For example, enhanced perception of visual detail, hypersensitivity to change, difficulties with facial recognition, and a tendency to play with parts of toys as opposed to the whole toy. These symptoms are sufficiently accounted for by WCC, as they can be explained by enhanced local perception and decreased global perception. For example, this local over global processing style can explain playing with parts of toys as opposed to whole toys (since there is greater awareness of the local level parts of the toy and lesser awareness of the global whole). Additionally, difficulties with facial recognition may also be explained by the fact that autistic individuals may not be processing the global details of the face that are required for successful recognition.

Furthermore, WCC could also potentially account for some of the social communication aspects of autism, such as difficulties with language learning, as language learning requires integrating what is being said with the relevant object or target in order to form an association.

Section 2.3.1.2 – Weaknesses of Weak Central Coherence Theory

The original account of WCC, which claims that there is a global processing deficit, has been challenged by empirical evidence that suggests that, although there is a default preference

towards local-level processing, global-level processing is intact in autism (Koldewyn et al., 2013). This finding has led to the development of a revised account of WCC that suggests that the local-over-global processing bias occurs due to a preference for local processing, as opposed to an inability to process globally (Happé, 2013 b).

However, this revised account still faces larger overarching problems about the claim that there is a local-over-global processing bias in autism populations. The concepts of ‘local’ and ‘global’ raise many valid philosophical questions and criticisms. This is due to the fact that ‘local’ and ‘global’ seem to be contextual terms. It is possible that what is considered to be a ‘local’ level stimulus in one scenario may be a ‘global’ level stimulus in another. For example, a lightbulb could be considered to be a constituent (local level) part of a lamp but may also be considered as an independent global whole in itself, as a lightbulb consists of many smaller constituent parts, such as the filament, fuse, and screw cap. This raises the further issue of whether anything can really be considered to be exclusively local or global. Is there a definitive point at which something can no longer be either: a) broken down into smaller constituent parts or b) a constituent part of something larger? These issues pose a significant threat to WCC, which is dependent on the idea that the essence of autism is a local-over-global processing bias. This is because even if there were something that was exclusively local and could not be broken down into constituent parts (perhaps something at the quantum level), it is clear that this is not what autistic individuals are perceiving. It seems that autistic individuals utilise global processing at least to some degree, as they are able to identify features of the environment that are global in the sense that they consist of smaller constituent local parts (such as lightbulbs).

While it might be straightforward to identify what is ‘local’ and what is ‘global’ in the controlled setting of specific studies (such as Navon tasks), it is significantly more difficult to characterise what is ‘local’ and what is ‘global’ in everyday settings. This is problematic for WCC, which attempts to use these terms to explain aspects of autism that are related to daily living (outwith such controlled experimental settings). A clearer explanation of the nature of local and global processing must first be provided in order for WCC to provide a sufficient characterisation of perceptual processing in autism.

Additionally, WCC is unable to account for many of the symptoms and behaviours that are common to autism. One major category of behaviours that WCC struggles to account for is RRBIs. It is unclear how restricted repetitive behaviours and interests are accounted for by

differences in global integration. For example, WCC struggles to account for special interests, stimming behaviours, and drive for repetition and difficulty with disengagement. In addition to this, there are some sensory aspects of autism such as sensory seeking and meltdowns in response to sensory overwhelm, which are difficult for WCC to explain. WCC also struggles to account for some differences in language utilisation such as echolalia. Furthermore, Cushing (2012) notes that WCC also struggles to account for differences in motor skills.

Section 2.3.2 – Enhanced Perceptual Functioning Theory

Similarly to WCC, Enhanced Perceptual Functioning (EPF) claims that the central element of autism is enhanced perception of local details (Mottron et al., 2006). However, the main difference between these two accounts is that WCC implies that enhanced local detail perception is the result of a trade-off between local and global processing where global processing must be impaired (classic WCC) or underutilised (revised WCC), while EPF claims that there is no such trade-off between local and global processes (Mottron et al., 2006). According to EPF, autistic individuals are less likely to rely on perception of global wholes as they have superior low-level local processing abilities. It is important to note that global processing is intact on this account (Corbett et al., 2016).

Section 2.3.2.1 – Strengths of Enhanced Perceptual Functioning Theory

EPF provides an account of autism in which the condition is positively framed. For many, this is a welcome difference from the typical deficit-based models of autism.

Additionally, EPF is well equipped to explain sensory sensitivities and improved detail perception in autism. If the essence of autism is enhanced perceptual abilities, then these aspects of the condition are to be expected.

Section 2.3.2.2 – Weaknesses of Enhanced Perceptual Functioning Theory

However, EPF is not able to account for all of the sensory differences in autism. For example, it only accounts for hyper-sensitivity and not hypo-sensitivity, and it does not account for difficulties with global processing or difficulties with multisensory processing.

Additionally, EPF cannot account for many of the non-perceptual aspects of autism. For example, it cannot account for any of the social communication differences in autism or any

of the symptoms or behaviours that are related to RRBIs. These differences are not clearly related to enhanced detail perception abilities.

Furthermore, the criticisms of WCC, which I outlined in section 2.3.1.2, (in relation to defining ‘local’ and ‘global’), will also apply to the EPF account since it also relies on the notion of local processing differences in autism.

Section 2.3.3 – Theory of Mind Theory

Theory of Mind (ToM) theory identifies the core essence of autism as a deficit in ‘mentalising’ ability, which is the ability to attribute mental states, such as beliefs, desires, intentions, or emotions, to other people (Baron-Cohen, 2000). Basic mentalising abilities are thought to develop by the age of 4 in neurotypical children; however, it has been argued that the development of these abilities in autistic children is significantly delayed, developing at around 9 years of age (Baron-Cohen, 2000).

ToM is supported by false belief tasks; these are tasks in which the participant must assign false beliefs to another person or character. The main test that is used is the Sally-Anne task, in which there is a scenario involving two dolls, Sally and Anne. The participant must observe where the first doll places their marble (either in the box or the basket within the scene) before they leave the scene. Then the second doll will enter the scene and move the marble to the other location, so that its placement has been changed. The first doll will then return to the scene and attempt to look for the marble. The participant is tasked with answering where the first doll will go to look for her marble. The expected, correct, answer is that she will look for it in the location where she left it (Carlsson et al., 2018). However, many autistic children under the age of 9, and neurotypical children under the age of 4 (Baron-Cohen, 2000), will incorrectly answer that she will look for it in the new location. These results are intended to show that those who fail the test do so because they lack mentalising abilities; they lack the ability to attribute the first doll with her own distinct mental states. Instead, they attribute their own mental states to her (by assuming that she knows where the marble really is because they themselves know where the marble really is).

Section 2.3.3.1 – Strengths of Theory of Mind Theory

The ToM theory accounts well for symptoms of autism that are related to social communication. For example, it accounts for reduced utilisation of joint attention and reduced use of protodeclarative pointing while still typically using protoimperative pointing (Baron-Cohen, 1989). If one does not consider others to be minded, then prompting others to share in joint attention would be pointless (since they would be assumed to lack the mental processes that are necessary for doing so). However, there would still be purpose in using gestures to request things, since this results in bringing about desired outcomes (regardless of whether or not others are considered to be minded).

This may also explain other social communication differences in autism, such as limited utilisation of or complete lack of speech, less engagement in two-way conversations, extreme directness and honesty, preferring to talk about their own interests and asking less about the interests of others, and atypical emotional responses to the situations of others. Without an understanding of the mental states of others, there is no logical motivation for utilising typical communication styles. For example, there is less motivation to consider how directness of speech may affect the mental states of others and less motivation to engage in conversations with others more generally, if the individual does not have an understanding that these other people have distinct mental states.

Section 2.3.3.2 – Weaknesses of Theory of Mind Theory

However, there are many other social symptoms and behaviours that ToM fails to account for.

ToM fails to account for echolalia, idiosyncratic language use, speaking in an unusual accent, reduced utilisation of eye contact, mismatching non-verbal and verbal communication, and lack of awareness. These are social symptoms of autism that are not explained by deficits in mentalising abilities.

ToM also does not account for any of the symptoms and behaviours related to RRBIs. For example, it cannot account for stimming behaviours, special interests, preference for sameness/aversion to change, or repetitive interests with difficulties with disengagement. Nor does it account for sensory differences within autism such as hyper- and hypo- sensory sensitivities and differences in multisensory processing. It also does not account for meltdowns, which may sometimes occur in relation to sensory overwhelm.

Additionally, masking seems to be counterintuitive on this account. It seems that in order for an autistic person to mask, in attempts to blend in more effectively, they must have an understanding that others have mental states that can be different from their own mental states. Masking is an attempt to change the mental states of others, namely the impression that others have of you and your behaviours, in order to fit in with their expectations. This requires knowing that others have expectations that are different from your own and knowing roughly what those expectations might look like in order to attempt to match them.

Furthermore, not all autistic individuals fail false belief tasks. In the original study by Baron-Cohen, Leslie, and Frith (1985), 20% of autistic participants did not fail the false belief task. If these tasks measure mentalising abilities, as ToM theorists suggest, then it seems that some autistic individuals are able to utilise mentalising abilities. This goes against the essence (a deficit in mentalising abilities) that is identified by ToM.

Additionally, these false belief tasks rely on other skills and abilities besides mentalising abilities; this means that failing a false belief task does not necessarily indicate deficits in mentalising. For example, language skills are important for the successful completion of certain false belief tasks (Polyanskaya et al., 2021). When the format of the task is changed, autistic individuals are more likely to pass these tasks (showing that the format of the task may be the source of the difficulty as opposed to mentalising) (Glenwright et al., 2021).

Moreover, autistic adults are typically able to pass standard versions of false belief tasks (Sommer et al., 2018). This suggests that if mentalising deficits are the core element of autism, then autism is a condition that could be outgrown.

Another criticism of the ToM account comes from Cushing (2012), who claims that many autistic children show signs of autism before the age at which any child would develop mentalising abilities (before the age of 4). This undermines the idea that deficits in mentalising abilities are the essence of autism, as differences between autistic and non-autistic children are observed before children from either group would typically develop mentalising abilities.

Furthermore, the double empathy problem provides evidence against the core claims of the ToM theory. It suggests that autistic people may have a different communication style as opposed to deficits in mentalising abilities. The double empathy problem highlights the fact that mentalising is also difficult for non-autistic individuals when communicating with autistic individuals (due to incompatible communication styles). Additionally, autistic

individuals tend to be able to communicate with, and understand the mental states of, other autistic people more easily than they can with neurotypical people (García, 2026). This shows that there is not a blanket deficit in communication or mentalising abilities in autism and instead suggests that these difficulties arise as the result of both groups utilising different communication styles, which may make it difficult for them to relate to each other.

Finally, the characterisation of autism on the ToM theory is potentially offensive. The ToM account can be taken to suggest that autistic individuals lack empathy or the ability to understand other people as minded. This clashes with the self-reported experiences of autistic individuals, who often report high levels of affective empathy even if they may not always have a cognitive understanding of the perspectives of others (Shalev et al., 2022).

Section 2.3.4 – Executive Dysfunction Theory

Cristofori, Cohen-Zimmerman, and Grafman (2019) define executive functioning abilities as follows: “Executive functions (EFs) include high-order cognitive abilities such as working memory, inhibitory control, cognitive flexibility, planning, reasoning, and problem solving. EFs enable humans to achieve goals, adapt to novel everyday life situations, and manage social interactions.” (Cristofori, Cohen-Zimmerman, and Grafman, 2019, p.197).

The Executive Dysfunction theory claims that the essence of autism is deficits in these executive functioning abilities (Hill, 2004).

Section 2.3.4.1 – Strengths of Executive Dysfunction Theory

It is because executive functioning is responsible for inhibitory control and behavioural flexibility that the Executive Dysfunction theory can account for many of the symptoms and behaviours of autism that relate to RRBIs. For example, it can account for insistence on sameness/aversion to change, repetitive actions (including stimming behaviours), and fixation on special interests (Hill, 2004).

These symptoms are intuitively explained by this account since if the core essence of autism is a deficit in executive functioning abilities, then these symptoms (which result from reduced inhibitory control and inflexibility) should be expected. Difficulties with inhibitory control could make appealing behaviours (such as stimming and fixation on special interests)

difficult to resist, while inflexibility could explain preference for sameness and repetition of these behaviours.

Section 2.3.4.2 – Weaknesses of Executive Dysfunction Theory

However, Executive Dysfunction theory cannot account for most of the other symptoms and behaviours that are common to autism. For example, Executive Dysfunction theory cannot account for perceptual differences such as hyper-sensitivity/hypo-sensitivity or reduced utilisation of multisensory processing.

It also cannot account for many of the social communication differences that are common to autism. For example, it cannot account for differences in the utilisation of language (including idiosyncratic language use), difficulty identifying emotions, reduced utilisation of eye contact, or mismatching non-verbal and verbal communication.

Although it may potentially account for some communication differences – such as talking to others about one's own interests exclusively, as this could be related to reduced inhibitory control.

Section 2.3.5 – Section Conclusion

All of the essentialist accounts that I have outlined in this section have failed to overcome the heterogeneity problem. They have failed to identify an essence of autism that is capable of unifying the wide range of symptoms and behaviours. Additionally, none of these accounts have even attempted to address the second prong of the heterogeneity problem – individual variance.

Section 2.4 - Eliminativism About Autism

The fact that none of these essentialist theories of autism have been able to successfully address the heterogeneity problem has led to an increase in eliminativism about autism (Anderson and Cushing, 2026).

Purported explanations of particular sets of supposed autism indicators like mindblindness, weak central coherence or executive dysfunction all failed to explain a sufficient number of the accepted indicators. Amongst those conducting the studies

and commentators thereon a consensus has emerged that, indeed, there is no essence of autism. (Anderson and Cushing, 2026, p.9)

Eliminativists believe that the term ‘autism’ should be abandoned as it does not pick out any real condition or group of people. It is instead an unhelpful and confusing label that wrongly groups people together. Throughout this thesis, I will be taking Cushing to be a representative proponent of autism eliminativism due to his clear criticisms against essentialist theories of autism and his firm commitment to this view (Cushing, 2012; Cushing, 2026).

Cushing (2012) claims that while there may be several real conditions relating to common symptoms of autism (such as sensory hyper-sensitivity), these conditions are wrongly grouped together under the label ‘autism’. He advocates for the conceptual elimination of autism.

He claims, as I have in the previous section, that essentialist theories of autism have all failed to overcome the heterogeneity problem and that, as a result, “these formerly trumpeted accounts of the essence of autism are steadily receding into the clinical past” (Cushing, 2026, p.208). Cushing (2026) explicitly references this failure of essentialist theories as a major motivation for eliminativism.

I drew an eliminativist conclusion from my anti-essentialism: without an essence, autism serves no purpose as a diagnosis. [...] ... we have diminishing reason to believe that there is something that the name autism corresponds to that is a real natural phenomenon. Instead, “autism” as a concept should fade away into obscurity, as have other terms associated with archaic inaccurate theories of the world, perhaps lingering as an inexact colloquial term—like “neurotic” or “hysterical” or “phlegmatic”—but having no place in psychology or indeed any of the sciences. (Cushing, 2026, p.208)

Cushing (2026) instructs his reader on how to best argue against eliminativism by stating that:

The best way to rebut my view would be to demonstrate that in fact autism is what is more generally referred to as a natural kind, that is, the groupings there in nature to be discovered. This could be done either by showing that it does in fact have an essence

(I won't hold my breath), or, alternatively, by arguing that it is possible to be a natural kind without having a clear essence... (Cushing, 2026, p.208)

Throughout this thesis, I intend to take the first approach – by providing a novel essentialist account of autism that can tackle the heterogeneity problem more effectively.

In the next chapter, I will examine a further essentialist account of autism—Monotropism — that has more potential to address the heterogeneity problem than the accounts that have been outlined in this chapter. However, I will ultimately argue that Monotropism is insufficient in its current form. Throughout the rest of this thesis, I will create a revised and improved account of Monotropism, which I have named 'Priority Monotropism'. In chapter 8, I will show how Priority Monotropism tackles the heterogeneity problem and how it can account for many of the features of autism that were outlined by the essentialist accounts in this chapter.

Section 2.5 – Chapter Conclusion

Throughout this chapter, I have shown that it is difficult to provide a clear definition of autism due to the heterogeneity problem. There is an extremely wide range of symptoms and behaviours that are common to autism. Additionally, these symptoms and behaviours are highly varied and difficult to conceptually link, as they cover social aspects, sensory aspects, cognitive aspects, and behavioural aspects of the condition. Another layer of the heterogeneity problem is added by the fact that there is individual variance between the symptom profiles of autistic individuals – this is to say that autistic individuals do not necessarily share the same set of symptoms or behaviours.

I have introduced several popular theories of autism: Weak Central Coherence theory, Enhanced Perceptual Functioning theory, Theory of Mind theory, and Executive Dysfunction theory. I have shown that while these theories are all able to account for some of the symptoms and behaviours of autism, none of these theories can account for the majority of these symptoms. Additionally, none of these theories can provide an explanation for the issue of individual variance in symptom profiles. Therefore, I have shown that none of these theories have been able to overcome the heterogeneity problem.

I have shown that the failure of these theories has led to increased adoption of eliminativism about autism. In the next chapter, I will introduce a further essentialist theory of autism that

has greater potential to address the heterogeneity problem—Monotropism. However, I will argue that it is insufficient in its current state. Throughout this thesis, I will develop a revised account of Monotropism in order to create a new theory of autism that is better able to address the heterogeneity problem and provide an alternative to eliminativism about autism.

Chapter 3 – A Further Essentialist Theory of Autism: Monotropism

Section 3.1 – Chapter Introduction

In the previous chapter, I have shown the difficulties that arise when attempting to develop an essentialist theory of autism. The main difficulty is the heterogeneity problem, which poses two challenges to an essentialist theory. The first challenge is identifying an essence that can account for the wide range of symptoms and behaviours in autism, and that can also explain how these seemingly unrelated symptoms and behaviours are connected. The second challenge is accounting for individual variance (accounting for the fact that autistic individuals may have different symptom profiles from one another). I have shown that the heterogeneity problem has been too great for the essentialist theories that I have outlined in chapter 2 to overcome.

In this chapter, I will introduce a further essentialist theory of autism: Monotropism. I will argue that Monotropism has the potential to account for a wider range of symptoms and behaviours in autism and hence has more potential to overcome the first prong of the heterogeneity problem. However, in this chapter I will argue that Monotropism is currently underdeveloped. I will argue that this is due to terminological, explanatory, and methodological limitations of the account. I will argue that many of these issues arise from ambiguities in the authors' use of the term 'attention' and the related term 'interest'. In the next chapter, I will consider the reasons why attention has not been clearly defined on this account. Throughout the rest of this thesis, I will work to create a revised account of Monotropism that provides a clearer definition of these terms and a clearer explanation of the role that attention plays in autism.

Section 3.2 – Introducing Monotropism

In order to characterise the extant account of Monotropism, I will utilise information from three key sources: Murray, Lesser, and Lawson (2005), Murray (2018), and Bogdashina (2016). Although Bogdashina is not one of the original authors, her work will be included in this discussion as she has made a significant contribution to the discussion of Monotropism, which has helped to clarify the perceptual aspects of the account.

I will first introduce the extant account of Monotropism in section 3.2.1. I will then argue that the nature of monotropic processing has been poorly characterised on this account. I will

argue that this is due to three main issues: terminological limitations, explanatory limitations, and methodological limitations.

Despite these limitations, there are positive features of the account. These features are based on two important elements: inclusion of the autistic community in autism research and substantial explanatory power for common symptoms and behaviours.

Due to the potential value of Monotropism, I will attempt to resolve these limitations. I will argue that these limitations largely arise as a result of the ambiguous characterisation of Monotropism. Therefore, I intend to attempt to resolve these limitations by providing a clearer interpretation of the claims that are made by the extant account, in order to create a more explicit characterisation of Monotropism. However, I will claim that ambiguities in the nature of attention cannot be resolved without the creation of a revised account of Monotropism that provides more explicit definitions and explanations about the nature of attention in autism.

Section 3.2.1 – The Extant Account of Monotropism

The extant account of Monotropism is outlined in a paper by Murray, Lesser, and Lawson (2005), which is titled ‘Attention, monotropism and the diagnostic criteria for autism’. In this paper, the authors acknowledge that autism is commonly diagnosed and characterised by a wide range of seemingly disconnected symptoms and behaviours, which can even appear to be in direct opposition to each other (as I have shown in my discussion of the heterogeneity problem in chapter 2). They state that Monotropism has the potential to show how these symptoms and behaviours are connected: these symptoms can be explained by an atypical distribution of attentional resources, which occurs during monotropic processing. On the extant account of Monotropism, this atypical distribution of attentional resources is the essence of autism.

This atypical distribution is a strong tendency towards monotropic processing, during which attentional resources are narrowly distributed to a select few interests while all other interests are unattended to. This process varies from the style of attentional distribution that is typically utilised by non-autistic individuals, which the extant Monotropism authors refer to as polytropic processing. During polytropic processing, attentional resources are widely distributed across a broader range of interests. To provide an example of the difference between these two styles of attentional distribution, Murray, Lesser, and Lawson (2005)

provide an analogy based on light sources. In this analogy, attentional resources are represented as light. Monotropic processing is represented as a torch beam – emitting a narrowly focused and intense beam of light, while polytropic processing is represented as a lantern – emitting a fainter, more diffuse glow of light.

Murray (2018) has since written a revised summary of Monotropism that further specifies the nature of monotropic processing. She states that “The central idea of monotropism [...] is that in autism, processing resource strongly tends to localize and concentrate to the exclusion of other input; an atypicality from which many other differences can be seen to follow.” (Murray, 2018, p.1). This explanation shows that, on this account, the narrow distribution of attention that occurs during monotropic processing is the essence of autism, which can explain a wide range of common symptoms and behaviours in autism.

Unlike the other essentialist theories of autism (which were discussed in chapter 2), the extant Monotropism authors have not conducted any empirical studies in order to support their theory; instead, they rely on the explanatory power of the account. There are two ways in which the explanatory power of the extant Monotropism is showcased. The first way in which the explanatory power of Monotropism is demonstrated is through its ability to explain a wide range of common symptoms and behaviours in autism. The second way in which this explanatory power is demonstrated is through the compatibility between Monotropism and self-reported experiences from autistic individuals. The extant authors were utilising a person-centred approach – the main goal of the theory was to account for the everyday experiences and behaviours of autistic individuals. However, as I will show throughout this chapter, this explanatory power is diminished by the lack of clarity about key terms and aspects of the account, as well as by the consideration of only a narrow range of self-reported experiences.

This person-centred approach was an important step forward for autism research, at the time of publication, in 2005. Many studies in this area were focused on ‘quantifying’ the autistic experience; relying solely on measurable results from empirical studies to support theories of autism was standard practice. Introspective accounts and the phenomenology of experiences were rarely considered by these types of accounts. This created a problem wherein many theories have attempted to characterise autism without considering the experiences of autistic individuals (as I have shown in my discussion of harmful theories in chapter 1). This is something that is still considered to be a problem in many areas of autism research today.

This focus on quantitative research has led many individuals in the autistic community to feel unheard, and a modern movement towards including autistic individuals and their subjective experiences in research about autism has garnered support (Loughran, 2020). Murray, Lesser, and Lawson (2005) recognised the importance of including autistic individuals in autism research long before this ideological paradigm shift. The recognition of this issue may have been due to the unique viewpoint that these authors had, as both autism researchers and autistic individuals. This meant they were sufficiently motivated to create an account that could bridge the gap between autism research and the autistic community. They did this by including autistic perspectives and by creating an account with qualitative foundations. This person-centred methodology demonstrated that the original authors were taking their subjects, and the phenomenological nature of their experiences, seriously.

Section 3.2.2 – Lack of Academic Recognition

However, while Monotropism has the advantage of considering and accounting for autistic experiences, the theory has not received as much academic recognition as many of the other essentialist theories that I previously outlined in chapter 2 (such as the Theory of Mind theory and the Weak Central Coherence theory). I will examine three main reasons why this might be the case by examining the following limitations of the account: terminological limitations, explanatory limitations, and methodological limitations. I will argue that many of these limitations are due to ambiguities in the account.

Section 3.2.3 – Terminological Limitations

The extant Monotropism account relies on a wide range of complicated and potentially ambiguous terminology. This is especially problematic as the authors have not sufficiently clarified these ambiguous terms. In this section, I will consider two of the most important terms that the extant account of Monotropism leaves unclarified: ‘interest’ and ‘attention’.

Section 3.2.3.1 – Ambiguities in ‘Interest’

One of the most important terms that is left unclarified is ‘interest’. This is particularly concerning as the authors characterise the mind as an ‘interest system’ this is a central feature of the extant account of Monotropism. The authors make an attempt to clarify this term by

stating that it is used in a way that ‘broadly coincides with common parlance’ (Murray, Lesser, and Lawson, 2005). However, I will argue that this explanation is insufficient, as the term ‘interest’ has many common parlance uses. It may be used to mean (1) ‘appeal’ as in ‘this subject is interesting’, (2) ‘attentiveness’ as in ‘this has captured my interest’, (3) ‘hobbies’ as in ‘fishing is an interest of mine’, (4) ‘concern’ as in ‘it is our best interest to take action’, or even (5) financial uses such as ‘interest rates’. While it is clear that definitions (4) and (5) are unlikely to be the intended forms of interest, it is unclear which of the other three definitions are being used. Definition (2) is the most relevant option for this account, as this is an attentional definition of interest. However, on the extant account, the term interest as is also used in ways that relate to (1) ‘appeal’ and (3) ‘hobbies’ (or personal interests). Throughout the extant account, the term ‘interest’ is used interchangeably between these three definitions without distinction.

However, it is important to note that this problem is not restricted to this account and may reflect a deeper discipline-level ambiguity. This interchangeable use of the term ‘interest’ may be due to symptom groupings used by the Diagnostic and Statistical Manual of Mental Disorders (American Psychiatric Association, 2013). As I have previously discussed in chapter 2, the DSM-5 lists Restricted, repetitive patterns of behaviour, interests, or activities as a category of autism symptoms. This category includes two separate sets of criteria that relate to personal interests/hobbies and perceptual interests/environmental attentiveness, respectively, while using the term ‘interest’ interchangeably between both categories without distinction. Murray, Lesser, and Lawson (2005) make significant reference to the DSM-4 (American Psychiatric Association, 2000) in this paper, which also uses the term ‘interest’ in this way. Therefore, the ambiguous use of the term ‘interest’ on the extant Monotropism account may have carried over from the DSM.

Section 3.2.3.2 – Ambiguities in ‘Attention’

However, Monotropism is faced with a further, more significant, overarching ambiguity problem – the extant authors have not clarified their use of the term ‘attention’. This is particularly concerning as Monotropism is an account that characterises the essence of autism as attentional differences. Attention is difficult to define due to its heterogeneous nature (I will discuss this problem further in chapter 4). Furthermore, the nature of attention is the subject of long-standing philosophical debate (as I will discuss in chapters 5, 6, and 7, where

I will discuss philosophical theories of attention). In this chapter, I will attempt to side-step the issue of characterising the nature of attention more broadly for now in order to focus on examining the core claims of the extant account of Monotropism about the nature of monotropic processing.

Section 3.2.4 – Explanatory Limitations

The essentialist theories of autism that were previously discussed in chapter 2 provide essences that are easy to summarise. These summaries help to make the accounts identifiable, memorable, and allow for the basic concept of their essence to be easily grasped.

Murray, Lesser, and Lawson (2005) do not provide a clear explanatory summary of this kind. The closest statement to a summary that could be formed from the extant account is the claim that ‘atypical distribution of attention is a central component of autism’. However, the nature of this atypical distribution is not clearly defined by this summary. This problem is partially resolved by Murray’s later account of Monotropism, in which she summarises Monotropism as follows: “The central idea of monotropism [...] is that in autism, processing resource strongly tends to localize and concentrate to the exclusion of other input; an atypicality from which many other differences can be seen to follow” (Murray, 2018, p.1). This summary is much clearer and provides some insight into what the nature of the atypical distribution of attention is like, for example, specifying that there is a tendency to distribute attention in a localised and exclusive style. However, this summary could be further improved. Here Murray (2018) is providing a claim about the theoretical account of Monotropism, namely that this account states that in autism there is a tendency towards a specific style of processing. This processing style is referred to, on the extant account, as ‘monotropic processing’. I will separate these terms and provide brief definitions.

Monotropism:

The theory that the central essence of autism is a difference in the distribution of attentional resources. More specifically, that there is a strong tendency towards monotropic processing in autism.

Monotropic processing:

A processing style in which attentional resources are narrowly distributed, focusing on a greatly reduced range of intensely experienced interests to the exclusion of many other potential interests.


In the following section, I will continue to provide a clearer, more explicit characterisation of monotropic processing by further interpreting the extant account.

Section 3.3 – Specifying the Character of Monotropic Processing

In the following subsections, I will specify the character of monotropic processing by identifying the following aspects of Monotropism: the monotropic/polytropic continuum, the attention tunnel, local over global processing, unimodal over multimodal processing, and style preference.

Section 3.3.1 – The Monotropic/Polytropic Continuum

Murray, Lesser, and Lawson (2005) state that there are two main styles of attentional distribution that exist in a continuum, with tightly focused attention at one end of the scale and widely spread attention at the other end. They use the analogy of light distribution, suggesting that at the tightly focused end of the scale attention is distributed like the beam of light from a torch (intense, narrow, and largely exclusive), while at the other end attention is distributed like light from a lantern (less intense, widespread, and largely inclusive). Monotropic processing is the attentional distribution style that occurs at the tightly focused end of the scale, while polytropic processing is the attentional distribution style that occurs at the widely spread end of the scale. Murray, Lesser, and Lawson (2005) propose that the “diagnosis of autism selects those few individuals at the deep or tight-focus extreme of this distribution of strategies” (Murray, Lesser, and Lawson, 2005, p.140). They suggest that autistic individuals are more likely to narrowly distribute their attention through the utilisation of monotropic processing, while non-autistic individuals are more likely to distribute their attention across a wider range of interests through the utilisation of polytropic processing. We can thus express this continuum as shown below in Figure 3.1.



(Figure 3.1)

Section 3.3.2 – The Attention Tunnel

In this section, I will specify that attention is distributed through an attention tunnel during monotropic processing. I will discuss the nature of this attention tunnel as it is characterised by the extant Monotropism account.

The narrow ‘torch beam’ of attentional awareness that is utilised during monotropic processing is referred to as an ‘attention tunnel’. The attention tunnel is highly exclusive, as only stimuli that are in focus within the attention tunnel are processed, while those that are outwith the attention tunnel are excluded. Bogdashina (2016) suggests that cutting a hole in a sheet of paper and overlaying this over an image while moving it around may help to portray visual perception in the attention tunnel.

However, while this suggestion helps to emphasise the extent of the perceptual exclusion during monotropic processing, the nature of the attention tunnel is left unclear. Since the extant authors have not commented on the nature of attention, it is difficult to consider how the attention tunnel should be understood. The extant account heavily implies a spatial account of attention—where ‘narrow’ attention within the attention tunnel may mean attention that is distributed to a narrow spatial area. However, the authors do not explicitly state that Monotropism is based on a spatial account of attention, and so it is not clear that the attention tunnel should be understood spatially. The attention tunnel may be another metaphoric representation of narrowly distributed attention. Furthermore, even if a spatial account is put forward, this only raises further unaddressed questions. For example, what size is the area of the attention tunnel? What shape is the attention tunnel? Is there any experience of stimuli that fall outwith the attention tunnel? Does the attention tunnel fade out into a

gradient or is there a sharp cut-off point? Is there only one attention tunnel or can multiple attention tunnels occur simultaneously? However, putting aside these ambiguities in attention, there are two key differences between monotropic processing and polytropic (typical) processing that are clearly expressed on the extant account. The first difference is a difference in the range of inclusion and exclusion, and the second difference is a difference in the intensity of the inclusion and exclusion. During monotropic processing, the attention tunnel is narrower, which means that the range of interests that are attended to is reduced. This narrower attention tunnel also means that there is a wider range of interests that are excluded. This claim holds regardless of whether attention is understood spatially or otherwise. Furthermore, Murray (2018) states that highly localised (narrowly distributed) attention leads to intensely experienced interests. However, again, it is unclear what is meant by ‘intensely experienced’; this may refer to the phenomenal character of the experience or to increased detail perception of the perceived interest, or perhaps both. Given the fact that phenomenally intense experiences and increased detail perception are both outlined by the personal experiences that have been used to support the extant account (which I will discuss in section 3.5), I will suggest that ‘intensely experienced’ should be understood as referring to both aspects.

Section 3.3.3 – Local Over Global Processing

As I have previously discussed in chapter 2, Weak Central Coherence theory (WCC) identifies the essence of autism as a local-over-global processing bias in autism that causes difficulties with integration and gestalt perception. Murray, Lesser, and Lawson (2005) acknowledge that Monotropism shares similarities with WCC. These accounts may be considered to be similar since both revolve around the idea of perception being ‘fragmented’ in autism. However, WCC accounts for this fragmented perception as occurring through a failure to sufficiently bind a wide range of attended stimuli together into coherent wholes (a failure of global processing). Monotropism accounts for this fragmented perception as occurring due to a narrower distribution of attention. This narrow distribution of attention means that a wide range of stimuli are unattended to, and so widespread inattentional blindness occurs. On the WCC account, autistic individuals could successfully attend to a wide range of stimuli (in a polytropic way), but they would fail to successfully integrate the attended stimuli into coherent wholes. Murray, Lesser, and Lawson (2005) emphasise this distinction by stating that “There may be no problems in integrating information *when it is*

attended to” (Murray, Lesser, and Lawson, 2005, p.141). This means that according to Monotropism, global processing is possible when all of the relevant information that must be integrated is attended to. This shows that narrowly distributed attention and local processing are distinct, since attention may be narrowly distributed and still include global wholes (Monotropism) or it may be widely distributed and exclude global wholes (WCC).

However, Murray, Lesser, and Lawson (2005) take this distinction a step further by stating that “from our perspective there is no reason to expect a preference for ‘local’ rather than ‘global’” (Murray, Lesser, and Lawson, 2005, p.142). However, I will argue that although there is no reason to expect a local-over-global processing preference, there is reason to expect a local-over-global processing bias on the extant account of Monotropism. In many cases, narrowly distributed attention may lead to the exclusion of global wholes. This is due to the fact that many of the local parts that make up these global wholes may be excluded via inattention, and as such full global wholes may not be sufficiently formed. This global exclusion may become more widespread with increasingly narrower ranges of attentional distribution. On this interpretation, the preference in autism is for monotropic processing, not local processing, but local-over-global processing may still be a consequence of monotropic processing due to the narrow attentional distribution that occurs during monotropic processing. Therefore, while the preference is not for local processing itself, it is for a processing style that may frequently result in local-over-global processing, and so a local processing bias may be expected in autism populations. However, it is important to note that this does not mean that Monotropism implies a global processing inability in autism, as global wholes may be attended to when sufficient attention is distributed to all of their constituent local-level parts, as is specified by Murray, Lesser, and Lawson (2005). It is simply the case that the narrowly distributed attentional style that is utilised during monotropic processing may frequently exclude many of these local-level constituent parts, and so global processing may be affected.

Section 3.3.4 – Unimodal Over Multimodal Processing

In the previous section, I have shown that the tendency towards monotropic processing in autism populations means that attentional resources are often narrowly distributed, which can lead to a local processing bias. In this section, I will discuss the idea that unimodal processing biases are also caused by this narrow distribution of attentional resources that occurs during

monotropic processing. This idea is discussed by Bogdashina (2016), who links monotropic processing to mono-channel (unimodal) processing. During multimodal processing, information from multiple sensory modalities is integrated into a coherent whole. For example, when watching television, the visual aspects are combined together with co-occurring auditory aspects in order to create a coherent audio-visual experience. During unimodal or mono-channel processing, information is only processed from one sensory channel at a time. For example, if mono-channel processing were to occur when watching television, then the viewer would only be able to process either the visual or the auditory elements; they would not be able to experience them together as a simultaneous audio-visual experience. Bogdashina (2016) characterises the narrow attentional distribution that occurs during monotropic processing as being intermodally narrow as well as intramodally narrow. On this view, the narrow attentional distribution in monotropic processing has the capacity to affect all sensory modalities and even cross-modal interaction.

Bogdashina (2016) states that mono-channel processing is a defensive strategy that is utilised in order to protect against sensory overload:

To avoid overload of sensory information, the brain consciously processes only one modality (although subconsciously a great amount of information may get in – ‘accumulation of unknown knowing’). The person might focus on one sense, for example, sight, and might see every minute detail of the object. However, while his/her vision is on, the person might lose awareness of any information coming through the other senses. Thus, while the person sees something, he/she does not hear anything and does not feel touch, etc. (Bogdashina, 2016, p.111)


This account suggests a style preference (as opposed to a deficit), as it shows that multimodal processing is possible in autism populations but that it may not be preferred since it can cause sensory overload.

Section 3.3.5 – Style Preference

Murray, Lesser, and Lawson (2005) state that the tendency towards monotropic processing results from a style preference as opposed to a deficit. While monotropic processing is the default style for autistic individuals on this account, it is important to note that autistic individuals are capable of utilising a wider attentional distribution style (since autism is not characterised by a deficit in polytropic processing).

They state that “There may be no problems in integrating information *when it is attended to*” (Murray, Lesser, and Lawson, 2005, p.141), showing that it is possible for autistic individuals to utilise global and multisensory processing. However, the process of widening the distribution of attention, in order to include a wider range of global/multimodal elements, consumes a greater number of attentional resources and requires additional effort for autistic individuals (Murray, 2018). On this account, the default processing style preference in autism is to narrowly distribute these limited attentional resources, and so global and multisensory elements are often disregarded.

According to Murray, Lesser, and Lawson (2005), the reverse is true of the non-autistic population, where there is a default preference to utilise polytropic processing. However, the authors are also not claiming that non-autistic individuals must always process in a polytropic way (with widely spread attention). It is simply that polytropic processing is the default style preference. On this account, the default autistic and non-autistic preferences are inverted. On this account, polytropic processing is the default style of attentional distribution for non-autistic populations, and monotropic processing is the default style of attentional distribution for autistic populations. Autistic individuals may override their default preference for monotropic processing and widen their area of attentional focus with effort, while non-autistic individuals may override their default preference for polytropic processing and narrow their area of attentional focus with effort (as is shown in Figure 3.2).



(Figure 3.2)

To summarise, the extant Monotropism account claims that there is a preference towards monotropic processing in autism populations. During monotropic processing, attention is distributed across a narrow range of interests, which means that many aspects of the environment are unattended to. This narrow distribution of attentional resources means that there may be local-over-global and unimodal-over-multimodal biases. However, it is important to note that Monotropism characterises the tendency towards monotropic processing as a processing style preference that can be overcome with increased effort.

In the next section, I will explain how these aspects of Monotropism may be used to explain a range of behaviours and symptoms that are common to autism, in order to show its potential for addressing the first prong of the heterogeneity problem (which was outlined in chapter 2).

Section 3.4 – Underexplained Symptoms and Behaviours

The authors of the extant Monotropism account claim that their theory can account for a wide range of common symptoms and behaviours in autism. Extant Monotropism is proposed to account for the following symptoms and behaviours: hyper- and hypo- sensitivities, co-ordination and integration difficulties, enhanced perceptual processing, difficulties with shifting attention, difficulties with contextual understanding, difficulties with mentalising, social inattentiveness, early language regression, difficulties keeping up with social cues, difficulties understanding social context, and experiences of overwhelm such as ‘meltdowns’ (Murray, 2018).

However, the extant authors underexplain how it can account for these symptoms and behaviours. The connection between monotropic processing and these symptoms and behaviours has not been made clear enough on the extant account. In this section, I will organise these symptoms and behaviours into three distinct categories: perceptual, cognitive, and social. I will then provide a clearer interpretation of how Monotropism may be used to explain these symptoms based on the core aspects of the account that I have previously outlined in this chapter. Throughout this discussion, I will also show how perceptual differences, which arise as a result of monotropic processing, can potentially explain many of the cognitive and social symptoms and behaviours. I will begin by examining these perceptual differences in the next section.

Section 3.4.1 – Perceptual Differences

In recent years, perceptual differences have been increasingly considered to be a central feature of autism. These perceptual differences can potentially explain other aspects of the condition, such as atypical behaviours, atypical cognition, and social difficulties (Schaaf, 2020). I will begin by explaining how the extant account of Monotropism could potentially account for perceptual symptoms and behaviours of autism, before going on to explain how these perceptual differences may give rise to cognitive and social differences.

In this section, I am categorising the following symptoms and behaviours, from Murray (2018), as being related to perception: hyper- and hypo- sensitivities, co-ordination and integration difficulties, enhanced perceptual processing, difficulties with shifting attention, and ‘meltdowns’.

Section 3.4.1.1 – Hyper- and Hypo- Sensitivities

Hyper- and hypo-sensitivities are commonly recognised symptoms of autism. A study by Marco et al. (2011) found that over 96% of autistic children reported experiencing hyper- and hypo-sensitivities. This suggests that hyper- and hypo-sensitivities are important symptoms to consider when identifying the essence of autism.

During instances of hyper-sensitivity, perceptual awareness is heightened to the point that sensory experiences can be uncomfortable or painful. This may lead the individual who is experiencing hyper-sensitivity to retreat to a location where sensory experiences are more controlled and comfortable. During instances of hypo-sensitivity, perceptual awareness is reduced. When experiencing sensory hypo-sensitivity, the individual may become understimulated and may seek out sensory stimulation in order to compensate. Hypo-sensitivities may explain an increased tendency to engage in restricted and repetitive behaviours in autism, such as spinning objects or eye rubbing, as these behaviours provide sensory stimulation.


The extant account of Monotropism explains the relationship between hyper- and hypo-sensitivities in two capacities: a causal capacity and a synchronous capacity.

Bogdashina (2016) describes the causal relationship between hyper- and hypo-sensitivities as follows: autistic individuals may experience broad hyper-sensitivities across multiple sensory modalities. This hypersensitivity causes sensory bombardment and leads to sensory overload.

In order to reduce this sensory overload, autistic individuals may utilise monotropic processing as a defensive compensatory mechanism to filter out the overload of incoming information and break environmental perception down into manageable components. This account suggests that hyper-sensitivities are the perceptual default in autism, while hypo-sensitivities are brought on by the utilisation of monotropic processing, which allows for the discomfort caused by hyper-sensitivities to be reduced.

However, on the extant account it is also the case that during monotropic processing hyper- and hypo-sensitivities occur synchronously. Murray, Lesser, and Lawson (2005) state that monotropic processing “results in large areas of potential information not being registered” (Murray, Lesser, and Lawson, 2005, p.146). This means that there is widespread hypo-sensitivity to sensory stimuli that fall outwith the narrow attention tunnel. However, Bogdashina states that monotropic processing can lead to the development of new, narrower, hyper-sensitivities within the attention tunnel. “As the person focuses on only one modality at a time, the sound may be experienced louder because it is all the person focuses on.” (Bogdashina, 2016, p.111). The difference is that these new hyper-sensitivities are more restricted and are more controllable; only stimuli that are actively attended to elicit hyper-sensitivities during monotropic processing. This removes the problem of widespread sensory bombardment.

This account addresses the puzzle of co-occurring hypo- and hyper-sensitivities in autism by explaining that both hyper- and hypo-sensitivities are to be expected since “there tends to be hyper-awareness within the attention tunnel, and [...] hypo-awareness, outside it” (Murray, Lesser, and Lawson, 2005, p.142). Furthermore, this account sheds light on a potential causal process suggesting that initial widespread hyper-sensitivities lead to sensory overload, which leads to the utilisation of monotropic processing as a compensatory strategy. This utilisation of monotropic processing then causes widespread hypo-sensitivity to stimuli outwith the current attention tunnel and narrow hyper-sensitivity to stimuli within the current attention tunnel (as shown in Figure 3.3).



(Figure 3.3)

However, there is still the issue of accounting for the timescale over which this process occurs. This process may occur over a long-term or short-term timescale. On a short-term timescale, the individual may wake up each day to face hyper-sensitivities and sensory overload before utilising monotropic processing. On this account, monotropic processing may be implemented on a situational basis, depending on whether the current sensory environment is particularly overwhelming. While on a long-term timescale, monotropic processing is utilised on a developmental basis due to a longstanding history of sensory overload. The extant account does not explicitly explain the timescale for this process. Evidence from the extant account supports a timescale that combines both situational and developmental timelines. The extant account supports the idea that a long-term preference for monotropic processing is developed but also supports the idea of short-term utilisation due to the fact that monotropic processing is not always utilised. On this account, a long-term developmental preference for monotropic processing is reinforced through the short-term instances where sensory bombardment occurs when monotropic processing is not being utilised. This supports the idea that there is a strong default tendency towards monotropic processing – although this can be overridden with increased effort on a situational basis.

Section 3.4.1.2 – Co-Ordination and Integration Difficulties

Co-ordination and integration difficulties in autism have led to the development of theories that focus on the idea of fragmented perception in autism (such as WCC, which was previously discussed in chapter 2). These accounts hold that autistic individuals may only experience select aspects of the environment and, as such, may have difficulty with perceiving whole scenes or successfully integrating sensory information from different channels. Monotropism shares this ideology but only while monotropic processing is active. This is an important distinction as, on the extant Monotropism account, autistic individuals

do not have a deficit in their ability to integrate sensory information. However, doing so will require overriding the monotropic processing preference and adopting a more polytropic processing style. Murray (2018) states that “Co-ordinating distinct interests (to forge a new whole) will require extra effort” (Murray, 2018, p.2). This shows that, on the extant Monotropism account, while integration is difficult and requires additional effort, it is not impossible.

Section 3.4.1.3 – Enhanced Perceptual Processing

As previously discussed in chapter 2, enhanced perceptual processing refers to the increased intensity of perceptual experiences and an increased ability to discriminate perceptual details in autism. Murray (2018) claims that “the more intensively localized [attention is] (at any given time) the more [intensely] experienced the current interest will be” (Murray, 2018, p.1). This suggests that enhanced perceptual processing is to be expected during monotropic processing, as attention is intensely localised through a narrow attention tunnel when monotropic processing is utilised. However, this enhanced perceptual processing should only be expected for stimuli that fall within the current attention tunnel.

Section 3.4.1.4 – Difficulties with Shifting Attention

Attention shifting refers to the ability to successfully redirect attention from one stimulus or set of stimuli to another or from one task to another. A visual example may be my ability to quickly look at the clock on my wall and then quickly return to look at my screen. An auditory example may be the ability to shift the focus of attention back and forth between a conversation and audio from a television in the room.

Murray, Lesser, and Lawson (2005) claim that shifting the object of attention requires a broad distribution of attention. Having a broad distribution of attention allows for an increased sense of awareness of the range of environmental objects between which the focus of attention may be switched. Conversely, having a narrow distribution of attention creates a decreased sense of awareness of environmental objects to which one could attend. Since attention is narrowly distributed during monotropic processing, autistic individuals may have difficulty with shifting their attention while they are utilising monotropic processing.

Furthermore, Murray (2018) claims that intensely experienced interests lead to the development of flow states. In these states, attention is fully absorbed by the current interest. Once in this flow state, the individual may have difficulty shifting their attention as their awareness of other stimuli is greatly reduced. Since monotropic processing includes intense experience of interests, autistic individuals may frequently experience this kind of attentional absorption in flow states and may frequently have difficulty with shifting their attention.

However, while this explanation shows why autistic individuals may have difficulty with switching attention (as it is more difficult for them to identify other stimuli in order to attend to them due to the narrow utilisation of attention), it does not sufficiently explain cases where attention does switch. Additionally, the nature of ‘flow states’ and what guides and instantiates them is unclear.

Section 3.4.1.5 – ‘Meltdowns’

Murray (2018) claims that disruptions to flow states can occur when an unexpected stimulus intrudes on the flow state. This is a negative experience that may cause distress and may result in the interrupted individual experiencing a ‘meltdown’. Murray, Lesser, and Lawson (2005) state that these interruptions result in a deeply negative experience.

To a person in an attention tunnel every unanticipated change is abrupt and is truly, if briefly catastrophic: a complete disconnection from a previous safe state, a plunge into a meaningless blizzard of sensations, a frightening experience which may occur many times in a single day. (Murray, Lesser, and Lawson, 2005, p.147)

On this account, these intrusions cause discomfort due to the fact that there are insufficient attentional resources left over to process the intruding stimulus. This is due to the fact that all of the attentional resources have been absorbed by the interest that was in focus during the flow state.

However, this explanation seems to suggest that during monotropic processing the autistic individual may still be highly aware of intruding stimuli. This is something that clashes with the previous core claims of the extant account about monotropic processing typically resulting in widespread unawareness. Additionally, it is unclear why this should be a negative experience for the individual – why should being pulled out of the attention tunnel cause a

“frightening” experience of plunging “into a meaningless blizzard of sensations”? This is something that is not sufficiently explained by the extant account.

Section 3.4.2 – Cognitive Differences

In the previous section, I have shown how perceptual differences may create a fragmentary experience of the environment. This fragmentary experience is caused by the fact that a limited range of environmental elements are attended to in great detail, while many contextual elements are disregarded. In this section, I will consider how these fragmentary perceptual experiences may lead to the development of idiosyncratic fragmented concepts that are formed on the basis of these perceptions.

In this section, I am categorising the following symptoms and behaviours, from Murray (2018), as ‘cognitive’: contextual understanding difficulties and difficulties with ‘mentalising’.

Section 3.4.2.1 – Difficulties with Contextual Understanding

Autistic individuals may have difficulties with understanding typical conceptual contexts or with connecting concepts together to form contexts. This difficulty is often referred to as ‘context blindness’ (Vermeulen, 2012). This difficulty with contextual understanding means that autistic individuals may frequently utilise concrete reasoning or fixed rules that do not vary based on the context of the current situation.

In individuals on the autism spectrum, interests will tend to be unconnected or idiosyncratically connected.” (Murray, Lesser, and Lawson, 2005, p.146). Fragmented perception, which is developed through monotropic processing, leads to the development of fragmented concepts – concepts which are detached from their original context. “... features of the environment which seem obvious to people with diffuse rather than tightly focused attention may be entirely missed. Monotropic individuals will tend not to recognize sequences of events, because no cognitive connection has been made between elements of the sequence. (Murray, Lesser, and Lawson, 2005, p.150)

These fragmented concepts may exist independently, or they may be linked together in a personal and atypical way to form atypical contexts. For example, seeing a hedgehog on a

day out to the beach may create an idiosyncratic connection between ‘beach’ and ‘hedgehog’. Additionally, reduced global attendance may mean that other common conceptual links are missed out on (‘towel’, ‘swimming’, ‘ice cream’, ‘seagull’, etc.). It is important to note that it is also possible for a typical conceptual link to be formed, such as ‘beach’ and ‘seagull’, when the typically associated stimulus is attended to. However, there is an increased possibility that atypical conceptual links will be formed since the range of stimuli that are attended to is significantly narrower.

Section 3.4.2.2 – Difficulties with ‘Mentalising’

As previously discussed in chapter 2, mentalising is the ability to attribute mental states to others and to interpret the mental states of others. For example, the ability to tell what emotions another person might be feeling or to understand what information they are likely to know. It has been suggested by Theory of Mind theorists that the essence of autism is a deficit in mentalising ability; this is supported by the results from false belief tasks (chapter 2).

However, (as I have already discussed in chapter 2) the results of studies using such tasks have not been unanimous. Murray, Lesser, and Lawson (2005) state that autistic individuals are able to pass false-belief tasks in lab settings once the task has been sufficiently explained and is understood. They argue that the real difficulty with mentalising ability stems from a reduced ability to successfully attend to all of the perceptual cues that are necessary for social understanding. Therefore, autistic individuals may pass mentalising tasks in controlled settings where the task is understood and there is a limited range of perceptual input.

However, they may struggle with these abilities in natural environments where a wider range of implicit perceptual cues must be attended to for a sufficient understanding of the situation. “In real-life situations, many crucial social cues occur very rapidly. Failure to notice them may lead to a general failure in assessing the meaning of entire situations” (Klin et al., 2003, p.345, cited in Murray, Lesser, and Lawson, 2005, p.145).

Section 3.4.3 – Social Differences

“We think the root of the social problems sometimes regarded as core in autism is probably attentional, and we are certain that those problems are worsened by the profundity of monotropism” (Murray, Lesser, and Lawson, 2005, p.144).

In the previous section, I have discussed the relationship between monotropic processing and the development of fragmented concepts. This relationship was defined as follows: monotropic processing leads to a narrow focus with widespread exclusion, this results in fragmentary perception on which fragmented or idiosyncratic concepts may be founded. These atypical concepts may lead to difficulties with social communication where meaning is ‘lost in translation’.

In this section, I will show that differences in both perceptual and cognitive styles may be responsible for the development of further differences in social styles. These social style differences may lead to difficulties engaging with others and with successfully conveying or interpreting meaning from social situations.

In this section, I am categorising the following symptoms and behaviours, from Murray (2018), as ‘social’: social inattentiveness, early language regression, keeping up with social cues, and difficulties with social context.

Section 3.4.3.1 – Social Inattentiveness

Autistic individuals may be mistakenly characterised as being ‘antisocial’ due to the fact that social inattention is a common characteristic of autism. As I have shown in chapter 2, autistic toddlers and children may not exhibit typical social behaviours, such as engaging socially or in play with others, responding to their name being called, or making eye contact during conversations. This may give the impression that the individual is unwilling to engage in social interactions with others.

However, the extant account of Monotropism suggests that this is not necessarily the case. Murray, Lesser, and Lawson (2005) state that

In a monotropic child, recognition of the existence of others will occur only in so far as other people are engaged with fulfilling the interests which pre-occupy that child. Otherwise, the existence of other people, like the existence of everything outside the tightly focused monotropic attention tunnel, may not impinge at all. (Murray, Lesser, and Lawson, 2005, p.148)

On this account, social inattention is caused by a reduced awareness of others due to the narrow distribution of attention during monotropic processing. Therefore, it is not necessarily

the case that autistic individuals do not want to engage in social interactions, but rather that they may have a reduced awareness of the social interactions that are going on around them, making it difficult to become involved.

Furthermore, fragmentary perception and cognition (developed through monotropic processing) may make it difficult for autistic individuals to understand social cues. Murray, Lesser, and Lawson (2005) state that: “On the phonetic level, sounds heard may not be identified as connected with each other but may be perceived as merely some among many noises in an unfiltered, undifferentiated aural environment” (Murray, Lesser, and Lawson, 2005, p.150). This may lead to a lack of understanding, even when social cues are noticed, since they may be perceived in a fragmented way.

Section 3.4.3.2 – Early Language Regression

As I have previously discussed in chapter 2, autistic infants may undergo a period of language development followed by a period of language regression. During this period of regression, autistic infants begin to disengage from using and responding to language. Interestingly, while this language regression is well documented, the cause is currently unknown (Al Backer, 2015).

Murray, Lesser, and Lawson (2005) state that Monotropism has the capacity to account for this language regression. They hypothesise that this regression is due to language being used as a method for capturing and controlling attention, especially during the critical language learning period in infancy, when this regression is often reported. They state that

such infants may start to learn speech as a way of expressing interest, and then be put off language by a change in how it is used in relating to them. As the infant's vocabulary gradually increases, other people start to use words as a way of seizing the infant's interest. For instance, a child may be looking at a ball but an adult may think that the child should be interested in the cat. [...] the adult points at the cat and says “cat!”. Once the infant has learned the word ‘cat’, the adult possesses a tool for manipulating the infants interest system. Disruption of the attention tunnel is a painful experience. Language may suddenly become unattractive for a deeply monotropic infant. (Murray, Lesser, and Lawson, 2005, p.151)

On this account, at this stage in infancy, language use becomes a tool for other individuals to manipulate or direct attention. This process causes a disruption to the toddler's current interest and requires their attentional awareness to be either shifted or expanded to include the suggested target. This process requires overriding the monotropic processing preference and expends additional effort, making this process uncomfortable. It is because of this discomfort that autistic toddlers may disengage with language use.

However, this also seems to clash with the previous claims of the extant account about monotropic processing leading to widespread unawareness. It is unclear how the adult could use language to seize the interest of the child if the child is not attending to the speech of the adult due to monotropic processing. A further explanation of how the speech of the adult intrudes on the fixation of the child is required.

Section 3.4.3.3 – Keeping Up with Social Cues

Social inattention may also be caused by difficulties in keeping up with social cues. A wide range of multimodal stimuli must often be attended to in order to successfully navigate social interactions (e.g. tone of voice, word choice, volume of voice, identity of voice – who is speaking, facial expressions, body language, directional cues, etc). During monotropic processing, there is a reduced awareness of simultaneously occurring stimuli, including multimodal stimuli. This is due to the fact that attentional resources are narrowly distributed to a reduced number of stimuli during monotropic processing. This means that many of the social cues that are necessary for successfully interpreting social situations may be missed. Therefore, individuals who utilise monotropic processing may have difficulties with successfully navigating social situations. Murray, Lesser, and Lawson (2005) emphasise this by stating that “The basic reason [that many autistic individuals may have difficulty with social participation] is the patchy and partial awareness that results from monotropic focus” (Murray, Lesser, and Lawson, 2005, p.148).

This difficulty in keeping up with social cues, which is caused by monotropic processing, may also mean that jointly attended social stimuli are not attended to. For example, if a group is looking together at a scene and discussing it, the autistic individual may miss out on the social cues that provide an invitation into joint attention (such as pointing or eye movements towards a point of interest) and as such they may continue to attend elsewhere and so will struggle to follow the discussion.

Furthermore, an increased tendency to enter into flow states during monotropic processing may cause further difficulties in attending to social cues, as attention may become absorbed by stimuli that are irrelevant to the current social situation. However, as I have previously stated, an improved explanation of what ‘flow states’ are and how they are instantiated and guided is necessary for them to be sufficiently explanatory.

Section 3.4.3.4 – Difficulties with Social Context

“Concepts about what might count for other people as a Gestalt may be absent” (Murray, Lesser, and Lawson, 2005, p.143).

The development of fragmented cognition may lead to the development of highly idiosyncratic concepts that are not accessible to others through a common language. For example, the development of an idiosyncratic concept of a ‘dog’ may lead to difficulties with engaging in meaningful conversation with others about dogs, as what the two speakers mean by the term is different. This may mean that the autistic individual may lose the meaning that was intended by the speaker or that the other individual may not understand the meaning that was intended by the autistic speaker, creating a communication barrier.

This is exemplified by Temple Grandin’s experience of identifying animals. She reported childhood experiences of fragmented perception and cognition in relation to categorising animals (Kawczynska, 2009). She claims that her perception is highly detail-based, focusing on individual aspects as opposed to wholes. Due to this detail-oriented style, she did not form conceptual Gestalts and instead relied on the identification of similar details in order to categorise animals. She had categorised animals on the basis of size, horses being large, dogs being medium-sized, and cats being small. However, this became problematic when she was introduced to different dog breeds that were the same size or smaller than cats. Grandin claims that she was able to resolve this problem through the realisation that all dogs have a similar style of nose regardless of their size. This example demonstrates her reliance on distinct details for forming concepts.

Section 3.4.4 – Recommendations for Practice

In this section, I will briefly review the recommendations for practice that are made by the authors of the extant Monotropism account. The following strategies may be used to promote social engagement, ease perceptual discomfort, and allow for effective learning.

Murray, Lesser, and Lawson (2005) suggest that those who are working with autistic children should ‘start where the child is’. This means that those working with autistic individuals should aim to engage with the individual’s current interest in order to promote learning, instead of trying to direct their attention away from the current interest. However, I will claim that this strategy would also be relevant for autistic adolescents and adults and, as such, it may be broadened to ‘start where the individual is’ or ‘start where the interest is’.

Murray, Lesser, and Lawson (2005) also suggest that task-irrelevant stimuli should be removed where possible, as this reduces the likelihood that the individual will become fixated on something that is irrelevant to the current task.

Additionally, Bogdashina (2016) claims that those working with autistic individuals should attempt to find out which channel is ‘open’ before trying to deliver meaningful information or requests. During monotropic processing, mono-channel processing may occur where information is processed exclusively through a single sensory modality at a time. If the individual is processing exclusively through vision and they are given an auditory instruction, then the individual is not going to be sufficiently aware of what they were asked to do.

Bogdashina (2016) recommends testing sensory channels to find out which one is actively being attended to in order to ensure that the individual is able to acknowledge the intended communication.

Section 3.5 – Self-reported Experiences

Murray, Lesser, and Lawson (2005) were focused on developing a theory that was inclusive of the experiences of autistic individuals. This is important as many of the studies of the time were less person-centred and focused largely on the results of studies or behavioural analysis without considering the personal experiences of the individuals who were under examination. Due to this inclusive person-centred methodology, Murray, Lesser, and Lawson (2005) included personal experiences in the form of self-reports from autistic individuals in their account. These self-reports exemplify the kind of experiences that occur during monotropic processing and give an introspective account of what it is like to utilise monotropic

processing. These self-reported experiences also relate to the range of symptoms and behaviours that were examined in the previous section.

However, Murray, Lesser, and Lawson (2005) do not provide an explicit interpretation of these self-reported experiences, nor do they clearly explain how these experiences exemplify monotropic processing or relate to the list of symptoms and behaviours that Monotropism accounts for. In this section, I will examine a selection of self-reported experiences from Murray, Lesser, and Lawson (2005) and Bogdashina (2016), which support the extant account of Monotropism. I will provide a brief interpretation of how these experiences relate to the core aspects of Monotropism and how they explain and/or exemplify the range of common symptoms and behaviours that were discussed in the previous section.

Self-reported experience 1

“[I was] intensely preoccupied with the movement of the spinning coin or lid, I saw nothing or heard nothing. People around me were transparent. And no sound intruded on my fixation. It was as if I were deaf” (Grandin and Scariano, 1986, p.20, cited in Murray, Lesser, and Lawson, 2005, p.141).

Here, Temple Grandin reports her own experiences of monotropic processing, showing that she was in a flow state where her attention was completely absorbed by a spinning coin or lid. As a result, she experienced social inattentiveness and blocked out auditory inputs, focusing unimodally on the local visual input of the coin or lid.

Self-reported experience 2

“I have many of the pieces of information stored within my memory and I can list them, but I cannot locate the connections that enable me to piece them together in a tangible way...” (Murray, Lesser, and Lawson, 2005, p.146).

This experience shows fragmented cognition in which concepts lack contextual backgrounds. This kind of fragmented cognition is to be expected as a result of monotropic processing, as utilising narrowly distributed attention leads to fragmented perception and widespread exclusion of surrounding contextual information.

Self-reported experience 3

I remember being around the table on a number of occasions. Sometimes it is to eat a meal, at other times it is to be part of some discussion group. These occasions can be very difficult because in order to comprehend well what is happening, I need to focus all of my attention on a number of activities all at once. For example, I need to look at what people's bodies are doing and at their facial expressions. I need to hear their words and process the whole event. [...]. After all of this I have to work out what my response should be. I often get this wrong because at times my attention is elsewhere focused and I miss the content and context of events... (Murray, Lesser, and Lawson, 2005, p.149)

This experience shows how monotropic processing can lead to difficulties in keeping up with social cues. This experience highlights the wide range of multimodal stimuli that must be attended to in order to successfully understand social situations. It is difficult to attend to all of the relevant elements when monotropic processing is being utilised, as it is likely that only a select few aspects will be picked up on. This leads to the development of a fragmented account of the social situation, which makes it difficult to choose an appropriate response.

Self-reported experience 4

I think that for many of us diagnosed as being on the spectrum of autism, the demand of having to 'pay attention' to so many things, simultaneously, is a nightmare. We tend to focus upon one thing at a time and this might mean we 'miss' lots of superficial information that gives context to much of life. (Murray, Lesser, and Lawson, 2005, p.152)

This experience showcases a general monotropic processing tendency and specifies that dividing attention, by utilising a polytropic processing style, is extremely difficult. This experience also shows that having a monotropic processing tendency may lead to decreased awareness of contextual information.

Self-reported experience 5

“My hearing became increasingly powerful... and I stopped seeing anything. I could focus all my concentration on only one sense, and that is hearing...” (Mukhopadhyay, 2008, p.6, cited in Bogdashina, 2016, p.112)

This experience shows unimodal processing with hyper-sensitivity to sensory input from within the ‘open’ sensory channel. This reinforces the idea that hyper-sensitivities may be experienced within the current attention tunnel.

Reported experience 6

Donna William’s inner-body sense, like everything else, was mostly in mono: if she touched her leg she would feel it on her hand or her leg, but not both at the same time. She had big restrictions being able to process information from the outside and the inside at the same time; for example, touching furniture, she could feel the texture of the wood but would have no sense of her own hand. She could also switch channels and feel her own hand but would lose sensation of what her hand was in contact with. (Bogdashina, 2016, p.112)

This experience shows the exclusionary nature of monotropic processing—attending to specific singular inputs and losing awareness of other unattended inputs. Importantly, this experience shows this kind of experience occurring locally at the intramodal level. This experience is also important as it shows that monotropic processing can apply intramodally to tactile sensations in the same way that it can apply intramodally to vision. This reinforces the idea that monotropic processing can apply to a wide range of sensory modalities.

Self-reported experience 7

Eye-contact is still difficult for me in noisy rooms because it interferes with hearing. It’s like my brain wiring lets only one sense function or the other but sometimes not both at the same time. In noisy rooms I have to concentrate on hearing. (Grandin, 2008, p.70, cited in Bogdashina, 2016, p.113)

This experience shows that autistic individuals may utilise monotropic processing as a compensatory mechanism to avoid sensory overload due to innate hyper-sensitivities.

Self-reported experience 8

This is like having one employee working in a huge department store instead of having a whole staff. If a customer has a demand in the shoe department whilst the worker is presently dealing with a demand in the toy department, the customer in the shoe department wouldn't find anybody there... and sometimes [this single worker closes] the shop in different departments when she can't cope with the demand. (Williams, 1996, p.99, 100, cited in Bogdashina, 2016, p.113)

This experience creates an analogy between the sensory network and a department store, where each sensory channel is represented as a different department. On this account, monotropic processing is similar to having a single member of staff attending to the whole department store. On this analogy, only one department (sensory channel) can be attended to at a time. The worker can generally move from department to department (attention can shift from one sensory modality to another), but when the worker is feeling overwhelmed, they may close some of the departments in order to reduce demand. This suggests that when an autistic individual experiences sensory overload, they may only attend to information from one sensory modality by closing off the others.

Self-reported experience 9

When I was very little, I remember forming wrong associations between words and objects. For instance, when I heard the word 'banana' while I was looking at a cloud, I labelled the cloud 'banana'. Then I'd get very confused when in another instance, I looked at the cloud and someone said the word 'table'. I would wonder whether some clouds were called bananas and some tables. (Mukhopadhyay, 2008, p.214, as cited in Bogdashina, 2016, p.130)

This experience shows how ‘misplaced’ narrowly distributed attention can lead to the formation of idiosyncratic concepts. Increased reliance on monotropic processing means that an autistic individual may often not be attending to the intended (‘typically expected’) stimulus and, as such, they may often form atypical associations. This experience shows how utilising monotropic processing can lead to difficulties with language learning.

Section 3.6 – Methodological Limitations

The ability to account for these common behaviours, symptoms, and experiences is a definite strength of Monotropism, which stems from its person-centred methodological approach. This allows Monotropism to potentially tackle the first prong of the heterogeneity problem more effectively than other essentialist theories of autism. However, I will argue that the extant account of Monotropism has some serious methodological limitations. The two limitations that I will discuss in this section are: a narrow range of self-reported experiences and a lack of empirical evidence.

Section 3.6.1 – Narrow Range of Self-Reported Experiences

While the extant Monotropism account has made important progress by including self-reported experiences from autistic individuals about the phenomenological nature of their experiences, this progress is diminished by the utilisation of a narrow range of self-reported experiences.

I have identified 16 self-reported statements that were presented in Murray, Lesser, and Lawson (2005). These experiences were provided by three different individuals: Donna Williams, an artist and writer who has written about her personal experiences of autism; Temple Grandin, a scientist and academic researcher who has also written about her personal experiences of autism; and finally, Wenn Lawson, an academic researcher and co-author of the paper in question. This range of self-reported experiences is very limited, with only three contributors, although this may be due to limited accessibility of self-reported statements at the time of writing. However, this limitation is further exaggerated by the fact that 14 of the 16 statements used in this paper are provided by Lawson, co-author of the paper. While I agree that Lawson has a unique and valuable position as both a researcher and as an autistic individual, this lack of variation creates serious problems for this account. Furthermore,

Lawson is personally invested in the development of the account, and therefore there could be a conflict of interest.

However, this problem is eased by Bogdashina (2016), who introduces a wider range of self-reported accounts that support Monotropism.

I will take this issue into consideration and ensure that I assess a wider range of self-reported experiences in the development of my revised account of Monotropism. Furthermore, I will not be relying on the personal experiences of Lawson due to this potential conflict of interest.

Section 3.6.2 – Lack of Empirical Support

Murray, Lesser, and Lawson (2005) aimed to provide a qualitative approach that would break away from the problems caused by the typical quantitative examination of autism. However, this approach has led to a lack of consideration of empirical evidence that may be used to support Monotropism. Murray, Lesser, and Lawson (2005) have engaged with results from empirical studies in order to dismiss other accounts such as Weak Central Coherence and Theory of Mind and to provide support for the idea that Monotropism is preferable over these accounts. However, the authors have not utilised empirical evidence in order to support Monotropism in general or to prove that monotropic processing is occurring in autism populations. While this qualitative approach is valuable and allows for an inclusive person-centred discussion, there is no reason that quantitative measures or empirical evidence should not also be utilised. There is no direct contention between these two approaches.

However, the problem is that, due to the ambiguous terms in this account, it is unclear how one would design empirical tests that could support extant Monotropism. While the authors of the extant Monotropism account have provided an idea for an essence of autism that has a lot of potential explanatory power, it has not been characterised clearly enough to realise that explanatory power. Critically, what is meant by important terms such as ‘attention’ and ‘interests’, as well as what it means for attention to be distributed in certain ‘narrow’ or ‘broad’ distribution styles, is left unexplained by this account. The lack of a clear explanation of how these terms are being used by this account both diminishes its explanatory power and prevents it from being empirically testable.

Section 3.7 – Chapter Conclusion

In this chapter, I have introduced a further essentialist theory of autism: Monotropism. I have provided an explanation of the core claims of the extant account of Monotropism and considered its potential explanatory power. However, I have shown that there are several serious limitations to this account. I have argued that the extant account does not provide a clear enough explanation of ‘attention’ or ‘interests’. This is something that has limited its explanatory power and empirical testability.

Throughout this thesis I will develop a revised account of Monotropism that will provide a clearer explanation of what is meant by ‘attention’ (and associated terms such as ‘interests’) in order to overcome the issues that the extant account faces.

In the next chapter, I will briefly examine the study of attention within psychology in order to gain an improved understanding of why the authors did not provide a definition of ‘attention’ on their account.

Chapter 4 – The Problem of Attention

Section 4.1 – Chapter Introduction

The purpose of this short linking chapter is to bridge chapters 3 and 5 and to provide an explanation for why I will be examining philosophical theories of attention (in chapters 5, 6, and 7) to provide a foundation for a revised account of Monotropism, as opposed to providing a definition of attention from psychology.

In this chapter, I will provide a brief recap of the problems that the extant Monotropism account faces in relation to the undefined use of the term ‘attention’. I will explain why it is especially problematic that the extant Monotropism authors did not provide a clear definition for their use of this term. I will then examine why the extant Monotropism authors did not provide a clear definition of ‘attention’. I will explain that this is likely the result of a deeper discipline-level problem within psychology, which is not exclusive to the extant Monotropism account. I will show that, like autism, attention also faces a heterogeneity problem that has led to eliminativist ideology and pessimism about the term.

Section 4.2 – The Extant Monotropism Account

As I have previously stated in chapter 3, the authors of the extant Monotropism account have not clearly defined their use of the term ‘attention’. This is especially problematic as this account identifies the essence of autism as attentional differences. As I have shown in chapter 3, on the extant Monotropism account, attentional differences are the core aspect of autism, and these attentional differences can explain many of the other symptoms, behaviours, and aspects of autism (such as perceptual, social, and developmental differences). According to this account, to be autistic is to have a preference for a narrow attentional style (monotropic processing). This is different from the way that attention is typically utilised (polytropic processing). However, without providing a clear definition of how the term ‘attention’ is to be understood, this statement is not providing clear claims about what autism is. What does it mean to utilise attention differently? What is it that is being utilised differently?

In chapter 3, it has been shown that there are many different types of attention that may be affected by Monotropic processing. For example, monotropic processing may affect perceptual attention (across a wide range of sensory modalities) or cognitive attention (across a wide range of cognitive processes). However, there are also many further types of attention

within these categories, such as overt/covert attention, endogenous/exogenous attention, and unimodal/multimodal attention, etc. The extant account does not explain which types of attention are affected by or are relevant to monotropic processing. Furthermore, the extant account does not specify how monotropic processing affects conscious awareness or perceptual experiences. However, the extant account does specify that such experiences are affected by monotropic processing and that this is a key difference between autistic and non-autistic individuals. These questions must be clearly answered in order for a sufficiently revised account of Monotropism to be developed.

There are only a few claims about the nature of attention that are clearly made by the extant Monotropism account. The first is that attention is a limited resource; this can be seen through the claim that “Attention is the resource which is competed for by task demand” (Murray, Lesser, and Lawson, 2005, p.141). The second claim is that there are (at least) two methods for the distribution of this resource: polytropic processing and monotropic processing. As I have shown in chapter 3, these two methods for attentional distribution exist as opposing poles of a continuum on this account. The third claim is that during monotropic processing attention is ‘narrowly’ distributed and during polytropic processing attention is ‘widely’ distributed. The fourth claim is that neurotypical individuals typically have a strong tendency to utilise polytropic processing while autistic individuals typically have a strong tendency to utilise monotropic processing.

However, the first and third claims seem to raise serious difficulties for this account since no definition for the term ‘attention’ has been provided. In terms of the first claim, the following questions arise: what is the resource that is limited? In what way is the resource limited? Why is the resource limited? Is this resource limited for both autistic and non-autistic individuals – or only for autistic individuals? Are there differences in the resource limitations between these groups? Are there different types of limitation for the different processing styles? The third claim also raises further questions: what exactly is meant by wide/narrow—what is it to attend to a ‘narrower range of interests’? Does wide/narrow refer to a numerical range of interests that can be attended to simultaneously? Or does wide/narrow refer to a specific spatial range? Can these terms be understood more broadly?

As I have previously discussed in chapter 3, one of the major problems with extant Monotropism is that it would be difficult to provide empirical support for the account. This is a substantial drawback of the extant Monotropism account, as other essentialist accounts of

autism, such as the Weak Central Coherence Theory, the Theory of Mind theory, and the Enhanced Perceptual Functioning theory, are at least arguably able to provide empirical evidence for their central claims about the essence of autism (as I have shown in chapter 2). I have argued that this difficulty arises from a failure to provide a clear definition of how the term ‘attention’ is being used on the extant Monotropism account. By not providing a clear definition of the term ‘attention’, the extant account has not made any explicit, testable claims about attention or how attentional processes vary between autistic and non-autistic individuals. While some important claims have been made, such as attention being a limited resource and attentional distribution being wide in polytropic processing and narrow in monotropic processing, these claims are too vague to be testable. In order to resolve this problem, we must first provide a clear definition for how the term ‘attention’ will be used and then explain how attentional resources and wide/narrow attentional distribution should be characterised under this definition.

In the next subsection, I will consider possible reasons why the authors of the extant Monotropism account did not provide a clear definition for their use of the term ‘attention’. I will show that this failure to define attention reflects a larger discipline-level problem within psychology.

Section 4.3 – The Discipline Level Problem

The following early psychological claims about attention showcase the potential importance of understanding attention: “... the doctrine of attention is the nerve of the whole psychological system...” (Titchener, 1908, cited in Sanders and Neumann, 1996, p.1), “The manifestations of the state that we commonly call attention are protean. No part of the individual is untouched by them. They extend to every part of the physical organism, and are among the most profound facts of mind” (Pillsbury, 1908, p.1, cited in Johnson and Proctor, 2004, p.2). From these statements, it is clear that attention is believed to be essential to the functioning of many mental processes.

However, despite the acknowledged importance of understanding attention, it remains a poorly defined term. The problem of not providing a clear definition for the term ‘attention’ is not exclusive to the extant Monotropism account and is part of a deeper discipline-level issue within psychology. It is extremely difficult to sufficiently define attention due to the fact that

it is highly heterogeneous (this will be further discussed in chapter 7). Hommel et al. (2019) demonstrate this heterogeneity by claiming that there has been:

... a dramatic variety of abilities attributed to attention: the ability to select external events for further internal processing (focused attention); ignore misleading information and/or an irrelevant location (selective attention); process irrelevant information (involuntary attention); selectively integrate information belonging to one event within and across sensory modalities (feature integration); prioritize processing of events from a particular location (spatial attention); systematically search for a target event (visual search); perform multiple tasks at the same time (divided attention); control the spatial parameters of eye movements (selective attention for action); prioritize one goal over others (goal-centered attention); prioritize one object, memory item, or conscious representation over others (object-centered attention); and consolidate information for later use and concentrate in anticipation of a possible event over some time (sustained attention). (Hommel et al., 2019, p.2290)

It is difficult to identify what, if anything, these widely varied types and instances of attention have in common, as no general neural correlate or mechanism of attention has been found (Hommel et al., 2019).

The above difficulties have led to attentional pessimism among attention theorists. Wayne Wu (2014) states that “Many contemporary theorists think dimly of the concept of attention, and this pessimism seems to be growing.” (Wu, 2014, p.4). Mole (2010) acknowledges similar concerns from Groos (1896), who claims that “To the question, “What is Attention?”, there is not only no generally recognized answer but the different attempts at a solution even diverge in the most disturbing manner.” (Groos, 1896¹, Mole 2010, p.4). While Sutherland (1998, cited in Wu, 2014) states that:

Over the past 50 years, the sheer ingenuity displayed by psychologists working on attention rivals if it does not exceed that of cosmologists studying black holes. Indeed, there is a similarity in their results – after many thousands of experiments, we know only marginally more about attention than about the interior of a black hole. (Sutherland, 1998, p.350, cited in Wu, 2014, p.4)

¹ Mole (2010) has cited Groos (1896) from Tsotsos, Itti, and Rees (2005) on page xxiv. However, Tsotsos, Itti, and Rees (2005) have not cited a page number for this quote from Groos (1896). Most authors are citing this quote through Tsotsos, Itti, and Rees (2005) despite this. However, I was able to find the following page number (p.210) cited in (Berlyne, 1974).

Furthermore, the following papers also demonstrate this issue of attentional pessimism:

‘There is no such thing as attention’ (Anderson, 2011), ‘Does attention exist?’ (Wilson, 2007), ‘No one knows what attention is’ (Hommel et al., 2019).

This growing attentional pessimism is especially concerning given the volume of research that has been conducted on the topic of attention. Dimension publication database (Digital Science, 2026) finds over 29 million publications on the topic of attention; this is over double the number of publications for the topic of perception (12.5 million) and substantially more than the number of publications for the topic of autism (around 900,000) (Digital Science, 2026). Despite this volume of research, it is acknowledged that “attention is certainly far from a clear or unified concept [which has] many, vague, and sometimes conflicting definitions...” (Lindsay, 2020, p.1).

Even William James (1890), who once famously stated that:

Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others... (James, 1890, p.403, cited in Wu, 2014, p.3-4)

later claimed that “Attention may have to go, like many a faculty once deemed essential, like many a verbal phantom, like many an idol of the tribe. It may be an excrescence on Psychology.” (James, 1890², cited in Anderson, 2011, p.1).

These issues of attentional heterogeneity and pessimism show that it would have been difficult for the authors of the extant Monotropism account to provide a clear definition of attention, since there is discipline-level difficulty with providing such a definition within psychology. In order to provide a clearer definition of how ‘attention’ should be understood for a revised account of Monotropism, I will examine philosophical theories that attempt to provide an answer to the metaphysical question: ‘what is attention?’. Over the next three chapters, I will introduce and assess the following three philosophical theories of attention:

² Anderson (2011) has not provided a page number for this quote from James (1890). However, I have found the following page number for it from a digital copy of the original text: (James, 1890, p.226).

the Selection for Action account (chapter 5), the Cognitive Unison account (chapter 6), and the Priority Structure account (chapter 7). I will provide both a general assessment of these theories and an assessment of their compatibility with the core claims of the extant Monotropism account. Once a suitable foundation for a revised account of Monotropism has been identified, I will outline the core claims of this revised account and further clarify its claims about the nature of attention during monotropic processing in chapter 8. This will allow for the development of an improved account of Monotropism that is better positioned to overcome the explanatory limitations of the extant account, which were previously outlined in chapter 3.

It is because of this discipline-level problem within psychology that I will examine philosophical theories of attention in order to specify how attention should be understood on a revised account of Monotropism. Psychology does not provide an answer to the metaphysical question: ‘what is attention?’. Instead, psychology examines certain types of attention and mechanisms that are associated with attention. This issue will be discussed further in chapters 6 and 7. In chapter 6, I will outline Mole’s explanation of why psychology has neglected the metaphysical question, and in chapter 7, I will provide Watzl’s arguments for why we should not reduce attention down to the mechanisms that are the focus of psychology and neuroscience.

Section 4.4 – Chapter Conclusion

In conclusion, in this chapter I have provided a potential explanation for why the authors of the extant Monotropism account have not clearly defined the term ‘attention’. I have shown that this is not a problem that is exclusive to their account, but instead it is a deeper discipline-level issue within psychology. I have briefly explained that attention also faces a heterogeneity problem, and I have shown the recent growth in pessimism about attention. In the next three chapters (chapters 5, 6, and 7), I will discuss and assess three major philosophical theories of attention and consider their suitability for providing a foundation for a revised account of Monotropism.

Chapter 5 – The Selection for Action Theory of Attention

Section 5.1 – Chapter Introduction

In this chapter, I will examine Wayne Wu's Selection for Action account. This is the first philosophical theory of attention, which will be discussed throughout this thesis, that attempts to provide an answer to the broader metaphysical question: 'what *is* attention?'. The primary text that I will be discussing in this chapter is Wu's 2014 book 'Attention'.

There are two main reasons for discussing Wu's account. The first reason is that Wu directly acknowledges the problem of attention and takes a philosophical approach to addressing this problem. Wu provides his own theory of attention, the Selection for Action account, which aims to resolve the problems that I have outlined in the previous chapter. The way in which Wu tackles the problem of attention is by introducing a framework for the construction of a general theory of attention. In this chapter, I will be examining both Wu's Selection for Action account and his framework, which may provide a useful foundation for the examination of other philosophical accounts of attention that will be discussed in chapters 6 and 7.

The second reason for examining Wu's Selection for Action theory is to assess its compatibility with the core claims of Monotropism in order to determine whether it would provide a suitable foundation for a revised account of Monotropism.

I will start by introducing Wu's framework and discussing his Selection for Action account in section 5.2, before going on to assess the suitability of this account for providing a foundation for a revised account of Monotropism in section 5.3.

Section 5.2 – Wu's Framework and the Selection for Action Account

Section 5.2.1 – Introducing Wu's Framework

Wu directly acknowledges the problem of attention that was outlined in the previous chapter. He claims that: "Many contemporary theorists think dimly of the concept of attention, and this pessimism seems to be growing." (Wu, 2014, p.4). Wu claims that this pessimism comes from the inability of any of the previously outlined psychological accounts of attention to unify the wide range of seemingly disparate aspects of attention. Interestingly, this problem

seems to mirror the heterogeneity problem in autism research, which was outlined in chapter 1. Wu acknowledges that most accounts of attention are able to provide insights into specific types of attention but are unable to provide insights into attention itself as a broader concept.

However, Wu argues that a successful account of attention needs to do more than provide us with a broad concept of attention. He claims that the Jamesian definition of attention has already provided us with a broad definition of attention by stating that “attention is selection for further processing” (Wu, 2014, p.13). This is a broad definition that would be compatible with the many different types of attention and is therefore a good contender for a broad account of attention. However, Wu (2014) claims that the Jamesian account of attention provides a definition that is too broad. He claims that the Jamesian account lets in too many counterexamples, which will be further discussed in section 5.2.2.

Instead, Wu (2014) claims that the way to resolve the problem of attention is to create a more specific account of attention as a broader phenomenon. He introduces a five-point framework that will allow broad accounts of attention to provide specific details by answering the following five questions:

The **metaphysical** question: what *is* attention?

The **function** question: what role does attention play?

The **properties** question: what are characteristic features of attention?

The **mechanism** question: How is attention implemented?

The **consciousness** question: what is the relation between attention and consciousness?

The following subsections, 5.2.2 – 5.2.6, will be dedicated to outlining Wu’s discussion of these framework questions. I will then provide a brief assessment of Wu’s methodology in subsection 5.2.8.

Section 5.2.2 – Question 1 – The Metaphysical Question

Wu (2014) claims that the first question that a successful account of attention must answer is the metaphysical question, which asks: ‘what *is* attention?’. I would argue that this is the main question that a theory of attention is trying to provide an answer to. This is the broad question that the extant Monotropism authors have failed to provide an answer to.

Wu (2014) states that his Selection for Action account provides a suitable answer to the metaphysical question for the following three reasons. First, because his account provides an answer to this question that avoids being too broad and therefore blocks unwanted counterexamples of attention, which are let in by the Jamesian account. Second, because his account is supported by current empirical methodology. Finally, because alternative accounts of attention are unable to provide a satisfactory answer to the metaphysical question.

I will begin by examining the first claim, that Wu’s account blocks unwanted counterexamples, which are let in by the Jamesian account of attention. As I have previously mentioned, Wu (2014) claims that the Jamesian account of attention is too broad and therefore lets in too many counterexamples of things that he argues should not be classified as attention. On the Jamesian account, attention is selection of information for further processing. Wu (2014) claims that the problem with the Jamesian account is that the notions of ‘selection’ and ‘further processing’ are too broad. He states that a successful theoretical account of attention needs to provide specific claims about the nature of ‘selection’ in order to remove problematic counterexamples.

Wu provides the following two counterexamples that he claims are problematic for the Jamesian account of attention. According to Wu, these counterexamples show instances where the Jamesian account would wrongly define a non-attentional process as attentional.

The first counterexample that Wu provides is computer processing. Wu claims that since computers are able to select information for further processing, computers could be said to utilise attention according to the Jamesian account. However, Wu argues that computer processing should not be considered to be attentional and therefore the Jamesian account has wrongly attributed attention to a non-attentional process (Center for Philosophy of Science, 2019).

The second counterexample that Wu provides is the process of colour spectrum filtering by the eyes. Wu claims that the eyes select visible light from the colour spectrum for further

processing and therefore the eyes could be said to be utilising attention on the Jamesian account of attention (Center for Philosophy of Science, 2019). Again, Wu argues that this is a process that should not be considered to be attentional and therefore the Jamesian account is, once again, wrongly attributing attention to non-attentional processes. He claims that the current Jamesian account must be revised in order to overcome these problematic counterexamples.

However, I will argue that these counterexamples are not as clearly problematic as Wu argues that they are for the following two reasons. First, it is unclear why Wu believes that these examples are counterexamples. He provides no reason why we should not consider these examples to be forms of attention, even if they are different forms of attention from those that are typically deployed by human beings at the subject level. It seems insufficient to dismiss the Jamesian account, based on the idea that it would be unusual to classify these examples as attentional, without providing sufficient reasons as to why it would be unusual to do so.

Why could it not be the case that these examples showcase specific types of non-human or non-subject attention, such as computer attention or eye attention? There are already examples of the concept of attention being used outwith the human/subject level. For example, the concept of ‘attention’ that is used in Artificial Intelligence systems (Soydaner, 2022) contends with Wu’s claim that the computer processing example should be considered as a counterexample.

Second, even if we accept that these examples should not be classified as instances of attention, I will argue that the second counterexample does not clearly pose a threat to the current Jamesian definition of attention. This is because it is not clear that the second example would be considered to be attentional based on the Jamesian definition, as it is not clear that the process of ‘selection’ is occurring in this example. It is arguable that colour spectrum filtering does not come about through a process of selection but rather through a physical incapacity. It seems mistaken to say that my eyes select the visible wavelength of light and select the filtration of UV and IR light. Rather, it seems to be the case that my eyes are physically incapable of detecting UV and IR light. There is no element of choice involved in this process. Human eyes do not choose not to detect UV or IR light, but are instead unable to do so. If this is the case, then this process does not involve selection and is therefore not providing a counterexample to the Jamesian account of attention.

In light of these counterexamples, Wu claims that the Jamesian definition of attention should be narrowed by further specifying the nature of selection. He then creates his own refinement of the Jamesian account by claiming that attention only selects information that is relevant for action (Wu, 2014). This forms the foundation of Wu's Selection for Action account of attention: attention is selection of information for further processing related to actions (Wu, 2014).

Wu claims that specifying the type of selection in this way will block the previously mentioned counterexamples that he takes to be problematic for the Jamesian account of attention.

However, I will object that the computer counterexample could still be applied to Wu's refined account. I will argue that specifying that the nature of the selection must necessarily relate to action does not allow Wu to avoid the computer processing counterexample. In fact, it may be argued that the primary function of a computer is to select information, on the basis of set rules, in order to perform actions. Therefore, a computer is selecting information for further processing related to actions.

I will note that Wu could argue that this counterexample is not applicable to his account, as this is not what is meant by selection. He may object by stating that selection is not a robotic processing of information in accordance with set rules and that selection is a more personal process that requires choice. However, if that is the case then this counterexample would also no longer apply to the Jamesian account of attention, as the computer would not be 'selecting' on that account either. This shows that the addition of processing being related to action is not the relevant clause that is blocking this counterexample. Therefore, this counterexample provides no sufficient reason to deviate from the current version of the Jamesian account.

Furthermore, while it is arguable that Wu's account still lets in 'counterexamples' of attention, it also seems to be the case that his account is too restricted and blocks out examples of things that should be considered to be instances of attention. The addition of 'selection for action' has made Wu's Selection for Action account very restricted. On this account, there can be no types of attention that are unrelated to action, as selection for action is a necessary component of attention. This blocks passive forms of attentional distribution that are unrelated to performing specific actions. For example, attending to the background hum of a refrigerator with no intention to act or automatically switching attention to a salient target (such as the sudden introduction of a bright colour) with no intention to act upon it.

Wu (2014) then continues on to provide a more specific refinement of his account, stating that: “where S perceptually selects X to inform performance of task T, then S perceptually attends to X (in T).” (Wu, 2014, p.48). He makes it clear that X can be a number of different strategies for the deployment of attention. This is a definite strength of his account, as this allows his theory to be flexible enough to account for many types of attention. For example, X may be a distinct spatial area or a specific feature. The nature of X may even change depending on the task that is being performed.

However, I have the following criticisms of this refinement. First, while this is not a definitive criticism, Wu’s formula seems unusual. It suggests that the subject attends to the mode of attentional distribution as opposed to the task or target. I will fill in the formula with the following two examples to demonstrate. Example 1: If subject (S) selects narrow spatial distribution (area) (X) to guide performance on completing a word search (T), then the subject attends to narrow spatial distribution (area) (X). Example 2: If subject (S) selects the colour feature red (X) to guide performance on the task of examining an apple (T), then the subject attends to the colour feature red (X). This suggests that objects are not attended to, but rather modes of attending to objects are attended to. This seems to go against our current everyday understanding of attention, where a subject would likely report that they were attending to the word search or the apple in the completion of the two example tasks above.

Second, Wu does not mention the underlying processes that allow for the selection of the specific attentional style (X) that would be best suited to the task at hand. This is important because (as I will show later in chapter 7) Wu criticises the Priority Structure account for including a process of prioritisation, yet it seems that a similar selection process is required for the subject to decide how X should be manifested. It seems that different attentional styles must be prioritised based on the task at hand.

The second argument that Wu (2014) gives, to support the idea that his Selection for Action account provides a suitable answer to the metaphysical question, is that his account is supported by current empirical methodology. Wu (2014) claims that it is because of this support from current empirical methodology that his account must be accepted in order for us to provide a sufficient answer to the mechanism question (which will be discussed later in section 5.2.5).

The reason that Wu claims that his account is supported by current empirical methodology is that his account is compatible with the way that we currently attempt to examine attention.

Current empirical methodology relies on the utilisation of tasks, which require participants to perform actions, in order to attempt to examine attention. Since these tasks involve participants performing actions, and action is a necessary part of Wu's Selection for Action account of attention, Wu (2014) claims that his account is well suited to our current empirical methodology. He claims that it is a benefit of his account that it allows us to resolve the problem of attention in a way that confirms that current empirical studies are actually identifying attention.

Furthermore, Wu (2014) claims that his account provides a suitable starting point for an attempt to define attention – since he claims that it is a known fact that attention is, at least sometimes, selection for task:

If theorists want to leave the mosh pit of attempts to state what attention is and not simply surrender, then the best option is to start with what everyone already knows in the experimental practice of attention: attention is, at least sometimes, selection for task. (Wu, 2014, p.271)

I am sympathetic to the idea that Wu's account provides a good starting place for the study of attention. I am also in agreement that a strong theory of attention should be able to assess the reliability of existing scientific research on attention. However, I will argue that while support from current empirical methodology could make Wu's account a good starting point for a theory of attention, this does not mean that his account is sufficient.

Additionally, it is not clear that his account does gain support from current empirical methodology. It seems unreasonable for an account of attention to attempt to gain support from current empirical methodology when one of the main purposes of developing an account of attention is to determine whether current empirical studies are actually measuring attention or not. Since we currently do not know what attention is, we cannot know if attention is what is being measured by our current empirical methodology.

If we were to accept Wu's Selection for Action account, then it would be the case that empirical studies are actually measuring attention, which he claims is a benefit of his account. It is understandable why this could be considered a benefit of the account, since it would bring about positive consequences such as affirming the results from a wide range of previously completed studies. However, Wu's account cannot be supported by current

empirical methodology if current empirical methodology is being supported by his account, as this support is cyclical.

Additionally, Wu (2014) claims that without accepting his answer to the metaphysical question, it would be impossible to provide a sufficient answer to the mechanism question. He claims that this is because attention needs to be related to tasks in order to establish evidence of underlying attentional mechanisms. This is because attention is not currently thought to be measurable without the use of tasks. Therefore, if we are to discover attentional mechanisms, then attention must be related to these tasks. Wu (2014) claims that "...in the search for basic mechanisms of attention, nothing is achieved without the empirical sufficient condition [attention is selection for action], a condition that gives neuroscientists grounds for concluding that the mechanisms and circuits that they uncover are "attentional"." (Wu, 2014, p.271). Since Wu states that a successful account of attention must answer the mechanism question he states that "...the action-centered account provides the best current answer to the metaphysical question: attention is selection for action." (Wu, 2014, p.271).

However, an immediate problem with this idea is that it is possible for attention to be measurable through tasks without actions being necessary for attention. We do not need to accept that attention is selection for action in order to provide an attentional account of actions. It might be the case that what attention *is* is not necessarily founded in actions but that actions are something that attention can affect. Examining actions performed during tasks could provide a way of measuring certain action-based correlates of attention. It is plausible that a mechanism of attention related to actions could be identified without holding that attention is always related to action.

The third reason that Wu (2014) claims his account should be accepted as the answer to the metaphysical question is that alternative accounts are unable to provide satisfactory answers to this question. Wu presents two branches of alternative views: accounts that hold that attention is selection for consciousness and accounts that hold that attention is selection for memory. He claims that neither of these accounts are able to give a satisfactory answer to the metaphysical question.

Wu (2014) refers to selection for consciousness accounts as 'gatekeeping' accounts in which attention serves as the gatekeeper to conscious experience. On these accounts, attention is required for conscious experience. He claims that these accounts do not provide a sufficient answer to the metaphysical question because they are unable to account for unconscious

forms of attention. However, it is important to note that in response to Wu's criticism, proponents of the selection for consciousness view may just claim that there are no unconscious forms of attention. However, this response would not allow selection for consciousness theorists to account for instances of what appear to be unconscious attention. For example, long-distance drivers attending to the road and the act of driving without having constant awareness of doing so.

Wu (2014) introduces selection for memory accounts as accounts that hold that attention serves the function of selecting information for memorisation. He claims that selection for memory theories are enveloped by his Selection for Action account. He claims that this is because selection for action often involves working memory processes and therefore selection for action will often also result in selection for working memory.

However, I will argue that these are not strong enough grounds to dismiss the selection for memory account as a standalone account. First, because working memory is not the only type of memory that selection for memory accounts may be concerned with. For example, these accounts could hold that attention selects for long-term memory, which would not be encompassed by Wu's Selection for Action account. Second, Wu claims that, on the Selection for Action account, attention *often* involves working memory processes. This is a separate claim from the stronger claim that selection for memory theorists are making, which is that attention *is* selection for memory and therefore always involves memory. Both selection for action theorists and selection for memory theorists would provide different answers to the metaphysical question. Selection for Action theorists would claim that attention *is* selection for action, while selection for memory theorists would claim that attention *is* selection for memory. This shows that while these two accounts may be related in some respects, they are distinct through their differing answers to this question.

Section 5.2.3 - Question 2 – The Function Question

The second question that Wu (2014) claims a successful account of attention should be able to provide an answer to is the function question, which asks: 'what is the functional role of attention?'.

He claims that a typical response based on the Jamesian model is that selection is the function of attention. On this account, attention is the process whereby some things are selected over others. Wu (2014) claims that this answer to the function question is too vague and that a

more specific claim needs to be made about the nature of this selective process. The main question that needs to be addressed is: ‘what type of selection is occurring?’. Wu attempts to answer this question with his Selection for Action account by stating that attention is the process whereby some things are selected over others in order to allow for action. He claims that the function of attention is to allow for the selection of one (or more) of many possible actions. This allows the agent to successfully perform actions without becoming overwhelmed by the many possible choices that they would have otherwise had.

Wu (2014) also makes other claims about the functional role of attention. He claims that attention plays a functional role in fixing demonstrative thought, enabling agency, determining and affecting the character of consciousness, allowing for epistemic justification, and fixing introspective thought. However, Wu has not clearly attached these claims to his Selection for Action account. Instead, these claims arise through the discussion of a wide range of contemporary theories of attention, each of which provides a case for one of these functions of attention. Wu states that these are important aspects of the functional role of attention and that a successful account of attention should provide an explanation that is able to account for all of these aspects.

Section 5.2.4 – Question 3 – The Properties Question

The third question that Wu (2014) claims a successful account of attention should be able to provide an answer to is the properties question, which asks: ‘what are the characteristic features of attention?’

Wu (2014) does not provide a clear answer to this question for his Selection for Action account, as he does not think that attention can be reduced to a single specific property. Instead, he emphasises that there is an extremely wide range of properties of attention that are recognised under current experimental paradigms.

He claims that “The science of attention will continue to uncover interesting features of this central psychological capacity, and as philosophers continue to explore the philosophical significance of attention, they will need to keep abreast of these developments.” (Wu, 2014, p.272). Showing that he believes that a sufficient philosophical theory of attention will need to take note of, and be able to account for, these different properties of attention as our understanding of attention develops.

Section 5.2.5 – Question 4 – The Mechanism Question

The fourth question that Wu (2014) claims a successful theory of attention should be able to provide an answer to is the mechanism question, which asks: ‘how is attention implemented?’

As previously discussed in section 5.2.2, Wu claims that acceptance of the Selection for Action account is necessary for a sufficient answer to be provided for this question. He claims that this is because action is necessarily involved in the tasks that are used to examine attentional processes and to look for potential attentional mechanisms based on our current empirical methodology.

However, it is important to note that while Wu (2014) claims that it is necessary to accept his account in order to answer this question, he does not identify a singular broad mechanism through which attention is implemented on his Selection for Action account of attention. Instead, Wu suggests that there are many different attentional mechanisms. He claims that a successful account of attention should be able to integrate these seemingly disparate attentional mechanisms that are present across the many different types of attention. However, he does not clearly explain how his Selection for Action account integrates these different attentional mechanisms.

Wu (2014) claims that further work is needed in order to move from a definition of attention to a full explanatory account that can bridge the gap between high-level theoretical aspects of attention and low-level mechanisms of attention. He states that:

A central task for cognitive science, in which philosophers will play a critical role, is not in defining attention, for we have the basis of such a definition, but in using that definition to integrate and bridge these disparate levels of analysis. Too often within cognitive science, work at different levels fails to be bridged in illuminating ways. That is hard work, and in the case of attention, work which presents interesting challenges and fertile ground for new approaches. (Wu, 2014, p.273)

Section 5.2.6 – Question 5 – The Consciousness Question

The fifth and final question that Wu (2014) claims a successful theory of attention should be able to provide an answer to is the consciousness question, which asks: ‘what is the relation between attention and consciousness?’

Wu (2014) argues that attention seems to be ‘closely tied’ to consciousness, but that the nature of this relationship is unclear. As such, he does not provide an explicit answer to the consciousness question for his Selection for Action account. Instead, he examines other accounts of attention and explains why their answers to the consciousness question are insufficient.

He claims that accounts that rely on conscious phenomenology in order to explain attention cannot provide a sufficient answer to the consciousness question. Wu examines a subgroup of selection for consciousness accounts, which he calls ‘phenomenological’ accounts. He describes these accounts as accounts that identify attention with conscious phenomenology. Wu (2014) argues that these accounts are insufficient, as he claims that attention cannot be reduced to a uniform phenomenology—since the phenomenological effects of attention are often disparate. He claims that a uniform phenomenology would be necessary for an explanation of this kind and therefore phenomenological accounts are unable to provide a sufficient answer to the consciousness question.

However, I have the following two criticisms of Wu’s objections to phenomenological accounts. First, Wu’s objection to phenomenological accounts only seems to hold if phenomenology is the sole answer to the metaphysical question. For example, an account may hold that phenomenology is an important aspect in answering the consciousness question but that considering phenomenology alone does not provide an answer to the metaphysical question. It may be the case, on an account of attention, that attention *is*, for example, selection for action, but that this account also allows that attention may typically have a certain phenomenology. Lack of uniformity would not be problematic for such an account because such an account could allow that different types of attention, for example, based on different types of action, could have different phenomenology. Such an account would not be locked into equating attention with a specific phenomenology.

Second, even if one is arguing that phenomenology is the sole answer to the metaphysical question, an important distinction must be drawn between the phenomenology of attention itself and phenomenological effects of attention. While the phenomenological effects of

attention may be clearly disparate, it is less clear that the phenomenology of attention itself is disparate. It may be argued that there is a quale of attention itself regardless of the content that is being attended to. This is to say that there may be a difference between the phenomenology of what is perceived through attending and the phenomenology of attending itself. Therefore, proponents of the phenomenological account could state that the phenomenology of attentional effects need not be unified so long as the phenomenology of attending itself is.

Wu (2014) also argues against the idea that inattentional blindness provides evidence for selection for consciousness or ‘gatekeeper’ accounts that state that attention is necessary for conscious experience. Inattentional blindness is the failure to notice aspects of the environment due to a lack of attention (Jensen et al., 2011). During instances of inattentional blindness, subjects report that they have not had any conscious experience of the stimuli that have gone unnoticed.

Wu (2014) makes it clear that, on his Selection for Action account, inattentional blindness does not occur, as he argues that there is no sufficient evidence for inattentional blindness. He provides the following two objections towards the concept of inattentional blindness. First, he provides criticisms from overflow theory. Overflow theorists believe that perception is possible without attention – on this view we are conscious of much more than we attend to. On this view, there is a separation between consciousness at the phenomenological level and consciousness at the access level. Access consciousness is consciousness that involves informational awareness, while phenomenal consciousness is consciousness that involves phenomenological experience. Wu (2014) sides with overflow theorists and uses overflow theory to argue that consciousness does not require attention. He claims that overflow theorists can account for evidence from inattentional blindness experiments by stating that the individual is conscious of the experience at the phenomenological level but not at the access level; this means that the subject has a conscious experience of which they are not aware.

Second, Wu (2014) provides criticisms about the nature of the ‘blindness’ aspect of inattentional blindness. He claims that it seems implausible that there is some kind of inattentional ‘hole’ or ‘void’ that is present during instances of inattention. This is based on the fact that he thinks that the phenomenology of such a void would be difficult to comprehend.

Instead, Wu (2014) claims that it is more plausible that something else is happening during instances of ‘inattentional blindness’. He claims that what may be happening instead during instances of supposed inattentional blindness is either inattentional agnosia or inattentional amnesia.

Wu (2014) describes inattentional agnosia as a state in which instances of inattention lead to a type of agnosia. During instances of inattentional agnosia, conscious experiences still occur and the properties of an object are detectable, but the subject may not be able to recognise the identity of the object to which these properties belong. This means that a subject who experiences inattentional agnosia may struggle to identify objects, which may lead them to fail to report experiencing the presence of said object even though they did have an experience of it.

For example, in the famous study by Simons and Chabris (1999), subjects were tasked with keeping track of the ball during a game of basketball; during this task, a man in a gorilla suit walks across the scene. This seems like it should be an obvious and salient stimulus that would be noticeable to the participant, but this study found that around half of the participants failed to notice the gorilla. Simons and Chabris use the results of this study to support the concept of inattentional blindness. However, according to Wu’s argument from inattentional agnosia, it could be the case that the subjects who failed to report the gorilla may have had a phenomenally conscious experience of the properties of the gorilla, (they may have been aware of its shape and colour for example), but they experienced agnosia towards the identity of the object – they did not register the object as a gorilla. Therefore, it did not register as anything particularly salient or surprising. On this account, the subjects did not report the gorilla because they did not know that the experience was of a gorilla.

I will argue that while this explanation may be able to account for higher-level (complex) stimuli that have an identity beyond basic features, this explanation cannot account for inattentional blindness regarding low-level (simple) stimuli, which may consist of very few features. For example, it seems unusual to suggest that a subject could have had a conscious experience of a red patch but experience agnosia towards its identity as a red patch.

Furthermore, this only seems to cover instances where there are a lot of distractors. For example, the properties of the gorilla may blend in with the properties of the background of the scene and with the other players who serve as distractors on this task. This may lead subjects to have difficulty with perceiving the properties of the gorilla as distinct properties of

a separate object or may lead them to wrongly attribute the properties of the gorilla to a background player instead. While it is the case that inattentional blindness experiments are often conducted in these types of busy environments, there are examples of instances where inattentional blindness occurs in environments that are relatively devoid of distractors. As I have shown in chapter 3, instances of inattentional blindness are often reported in autism populations even in low sensory environments where there are not a lot of distractors. For example, autistic children are often reported to fail to notice their name being called even in a quiet room.

However, it is important to note that I am not suggesting that inattentional agnosia is not possible or that it never occurs. I am simply arguing that there also seem to be cases of ‘inattentional blindness’ that are not covered by Wu’s suggestion of inattentional agnosia.

Wu (2014) claims that instances of what are referred to as ‘inattentional blindness’ may instead actually be instances of ‘inattentional amnesia’. He claims that the subject may have a conscious experience of unattended objects but may forget that they have had such an experience, as this experience was not stored in working memory (since the unattended objects were not task-relevant). This leads the subject to report that they were not aware of an object that they were actually aware of.

Wu (2014) claims that it is difficult to disprove the idea that instances of supposed inattentional blindness are actually instances of inattentional amnesia, since the subject would not have access to the relevant memories to report otherwise. I am sympathetic to the idea that it would be difficult to prove that subjects are not experiencing inattentional amnesia, since we are dependent on their reports, and they would have no awareness of the fact that they had forgotten.

However, this seems unusual for cases where the subject must give a live report of their experiences. In the case of live reporting, there would be significantly less time between the task and reporting stage for the participant to forget. A study of this kind has been conducted by Ward and Scholl (2014), where participants were tasked with live reporting their experiences without any time or need for the information to be stored in memory. This study found that participants still experienced inattentional blindness, and thus Ward and Scholl (2014) state that this provides evidence for inattentional blindness as a distinct phenomenon from inattentional amnesia.

Furthermore, a study by Bredemeier and Simons (2012) has shown that while working memory abilities affect performance on tasks with high attentional demands, working memory abilities do not affect the likelihood of participants experiencing inattention blindness on such tasks.

Additionally, I will argue against Wu's criticisms of inattention blindness based on the phenomenology of the blind spot being difficult to comprehend. First, I do not think that the phenomenology of a blind spot being 'odd' or difficult to comprehend is a sufficient criticism against inattention blindness. Second, it is not clear that inattention blindness necessarily creates such a 'void' or 'blind spot'. There are clear everyday examples of blind spots being filled in automatically. For example, the filling in of blind spots is part of our normal everyday perception, as each eye has a blind spot where there are no photoreceptors. However, this blind spot is not noticeable as it is automatically filled in (Durgin, Tripathy, and Levi, 1995; Raman and Sarkar, 2016). It is unclear why the same process of filling in the blind spot could not happen during instances of inattention blindness.

Wu (2014) provides similar criticisms against the concept of blindsight, although he does not argue that blindsight does not occur, but rather that attention is not necessarily present in blindsight. Blindsight is a perceptual condition in which individuals report being blind and therefore report having no visual experiences. However, in some cases, individuals with blindsight are able to make correct choices about the objects in their environment through vision alone. For example, an individual with blindsight may be able to visually confirm whether an object is colour A or colour B, shape A or shape B or in position A or position B during a forced choice task, despite reporting that they cannot see anything in their environment (Weiskrantz, 2007).

Wu (2014) discusses blindsight by claiming that attention and perception are separate states that are not necessarily related. He states that:

Psychological states, like perception, memory, and thought about an item, do not on their own conceptually entail attention to that item. [...] Similarly, unconscious perceptual states of an object, such as those in blindsight, do not entail attention to the object. Something further is needed if attention is to emerge from a subject's mental states. (Wu, 2014, p.91)

Wu states that selection for action is necessary in order for any of these states to become attentional states.

It is important to note that Wu is not claiming that attention is never present during blindsight. According to Wu's account, attention is utilised in cases of blindsight that involve selection for action. For example, on Wu's account, attention may be present in blindsight during the forced choice task examples above, since the participant may perform actions in service of these tasks.

Section 5.2.7 – Framework Overview

As discussed in section 5.2.1, Wu has only provided an explicit answer to two of the questions in his framework: the metaphysical question and the function question. His answer to the metaphysical question is that 'attention is selection for action' and his related answer to the function question is that 'the function of attention is to select for action'.

Wu (2014) optimistically claims that by providing this answer to the metaphysical question, he has shown that "We do know what attention is and this knowledge puts us in a position to investigate and discover what attention does, what it is good for, and how it works." Wu (2014, p.274). He claims that we do know what attention is based on his answer to the metaphysical question and states that accepting his answer to the metaphysical question allows us to be in a position to build an account of attention that examines the answers to the remaining questions.

Section 5.2.8 – Assessment of Wu's Methodology

Wu (2014) has outlined a clear framework of questions to which a successful account of attention should be able to provide answers. This is important as it helps to set a standard that future accounts of attention can strive to meet. Assessing existing accounts of attention by their answers to these questions will help to clarify where these accounts stand on major issues and distinguish accounts that may share central similarities. Wu has assessed many contemporary accounts of attention and has shown the importance of answering the questions in his framework. The development of this framework, where no such framework existed before, is an important contribution to the philosophical study of attention. Going forward, I

intend to use this framework in my examination of other philosophical theories of attention, in order to try to draw out as many explicit claims about the nature of attention as possible.

Throughout this section, I have criticised Wu's answers to the five questions that he has outlined in his framework. However, it is important to note that Wu seems to acknowledge that his account is incomplete, as he states several times that it provides a starting point for the examination of attention. However, I have argued against his claims that his account provides the best starting point and that other theories of attention are insufficient.

Section 5.3 – Compatibility with Monotropism

As discussed in the previous chapter, the main purpose for identifying a sufficient theory of attention is to provide a foundation for a revised account of Monotropism. In this section, I will examine the compatibility between Wu's Selection for Action account of attention and the core claims of Monotropism.

I will argue that Wu's Selection for Action account is incompatible with the core claims of Monotropism and therefore it would not provide a suitable foundation for a revised account of Monotropism. Wu's account is incompatible with Monotropism for the following reasons.

First, the concept of inattentional blindness is significant for the core claims of Monotropism. A large portion of the explanatory power of Monotropism comes from the fact that it is able to provide a reason for frequent widespread inattentional blindness in autism populations. This widespread inattentional blindness can then be used to explain further perceptual and social differences in autism, such as not responding to one's own name being called, difficulties in language development, sensory hyposensitivity, etc. Therefore, Wu's Selection for Action account would be unsuitable as a foundation for a revised account of Monotropism since Wu has made it clear that inattentional blindness does not occur on his account.

Second, Monotropism allows for instances of monotropic processing that are not necessarily related to action. For example, monotropic attentional styles can be utilised during instances of passive attention, such as listening to birds or looking at a chair with no intention to act upon either. Since Wu claims that attention necessarily involves action, these examples are not instances of attention on his Selection for Action account. Therefore, Wu's Selection for Action account would provide an unsuitable foundation for a revised account of

Monotropism as it would only be able to account for instances where monotropic processing is utilised during the performance of actions.

Section 5.4 – Chapter Conclusion

In conclusion, in this chapter I have introduced and discussed Wu's Selection for Action account of attention. I have provided a general assessment of his arguments and methodology, and I have argued against many of his criticisms that were directed at competing theories of attention. I have ultimately concluded that his Selection for Action account would not provide a suitable foundation for a revised account of Monotropism, as it would not allow such a revised account to account for instances of passive attention. However, it is important to note that while I will not be using Wu's Selection for Action theory as a foundation, I will be using his framework for the assessment of the other philosophical theories of attention, which I will discuss in the next two chapters.

Chapter 6 – The Cognitive Unison Theory of Attention

Section 6.1 – Chapter Introduction

In the previous chapter, I introduced and assessed Wu's Selection for Action theory of attention. In this chapter, I will introduce and assess Mole's Cognitive Unison theory of attention. The primary text that I will be discussing in this chapter is Mole's 2010 book 'Attention Is Cognitive Unison: An Essay in Philosophical Psychology'.

I will first outline the key parts of Mole's discussion concerning the problem of attention, in which he provides historical context for this problem. I will explain his adverbialist approach to answering the metaphysical question. I will then discuss and assess his criticisms of other theories of attention (including other versions of adverbialism). I will outline his Cognitive Unison theory of attention, on which to attend is to do something with 'cognitive unison'. I will provide an assessment of this theory before considering how it might be used to fill in Wu's framework (which was introduced in the previous chapter). Finally, I will consider the suitability of this theory for providing a foundation for a revised account of Monotropism. I will argue that it would provide an unsuitable foundation for such a revised account, as the way that Mole characterises attention in general is too close to the way that Monotropism characterises attention in autism.

Section 6.2 – The History of the Problem of Attention

I have previously outlined the problem of attention in chapter 4. In the development of his Cognitive Unison theory of attention, Mole (2010) further emphasises this problem and provides historical context that helps to explain how we arrived at such a problem.

Section 6.2.1 – William James

Mole is critical of William James's famous claim 'everyone knows what attention is...'. He argues that by making this statement, James has provided psychologists with an excuse not to further examine the fundamental nature of attention. He goes on even further to state that James's famous statement was a 'bluff', and that even at that time many other researchers working in the area were confused about what attention is.

James's remark that 'everyone knows what attention is' has an air of unanswerability about it, and an air of finality. It suggests that a more philosophically ambitious analysis of attention [...] would be superfluous. But this air of unanswerability should be looked on with suspicion. It is a signature trait of the bluffer. James's remark was, I think, a piece of bluff. (Mole, 2010, p.4)

While I fully agree that overreliance on James's famous claim has hindered the progress of developing a full understanding of attention, I feel that this is a little uncharitable. It may be the case that what James meant by this statement is that we all intuitively have some sort of common sense understanding of attention based on our experiences of it. It is not clear that his statement is about having a full conceptual understanding of the metaphysics of attention.

Section 6.2.2 – The Metaphysical Issue

Mole claims that the main problem in attention research is that the metaphysical question, what *is* attention?, has been left unaddressed for too long. This is something that has also been addressed by Wu as a core problem in attention research (as I have previously discussed in chapter 5).

He then goes on to explain the various historical factors that he claims have conspired against the development of our metaphysical understanding of attention.

Mole (2010) claims that at the start of the 20th century psychologists were keen to ignore the issue and that the rise of logical positivism within philosophy around this time further enabled this. During the first half of the 20th century, behaviorism dominated psychology. Behaviorism struggled to accommodate attention, as attention is not a publicly observable process (although it may sometimes produce publicly observable behaviours, it does not always necessarily do so). As a result, psychological research into attention was deprioritised at this time. The Second World War then brought attention back to the forefront of psychology as there was a sudden urgency to understand the limits of attention. Suddenly it was important to understand how many items on a radar screen an operator could reliably attend to and how to best direct and redirect the attention of others. Behaviorism was pushed out, and attention came back to the forefront of psychology. However, the study of attention was largely concerned with these practical issues, and this ultimately led to the metaphysical question being ignored. Attention researchers were largely occupied with examining the

utilisation and limitations of attention, as opposed to developing a metaphysical understanding of attention as a broad phenomenon.

Section 6.3 – The Metaphysical Question

Section 6.3.1 – Difficulties with Answering the Metaphysical Question

Mole (2010) claims that the metaphysical question is the first thing that needs to be answered in the construction of a theory of attention. However, he identifies several difficulties with providing an answer to this question.

He claims that the main issue with providing a metaphysical account of attention is that there are too many viable contenders for what attention is. This is to say that attention is too heterogeneous, as was previously discussed in chapter 4. It is difficult to create a theory of attention that can account for all instances of attention while still being narrow enough to exclude instances of non-attention and, in addition, can also say something meaningful about what attention is.

He claims that it is a fundamental mistake to attempt to answer the metaphysical question by cataloguing the constituent processes. This is because Mole favours an adverbialist approach, as opposed to a process-based approach.

Any process whatever, of the body or the mind, will be a case of employment [attention], as long as it is done in the right way and for the right reasons. A theory of [the] employment [of attention] should specify the way and the reasons. (Mole, 2010, p.13)

Section 6.3.2 – Addressing the Metaphysical Question

Mole (2010) claims that a successful theory of attention should allow us to answer both of the following questions:

First, it should allow us to provide an answer to the question: ‘what is attention?’

Second, it should allow us to provide an answer to the question: ‘what is it for something to be done attentively?’

Mole claims that there are two approaches that one can take towards answering these questions. The first approach is to answer the second question by answering the first question. On this approach, one would create an account that first identifies what attention is, and then explains that to do something attentively is to do the thing in the manner that typifies instances of this process of attention. Alternatively, one could take the second approach and answer the first question by answering the second question. On this approach, one would create an account that first identifies what it is for something to be done attentively, and then explains that attention is the doing of something in this attentive manner.

Mole (2010) claims that the first approach is typically used for explaining processes while the second approach is typically used for identifying phenomena. He provides the following two examples to demonstrate this: ‘combustion’ and ‘haste’. He explains that for processes, such as combustion, we would typically use the first approach (providing an answer to the first question before providing an answer to the second question). He states that we would first provide an explanation of what combustion is, and then we would be able to provide an explanation of what it is for something to happen ‘combustively’. He then states that for identifying phenomena we would need to use the opposite order of explanatory priority by taking the second approach – answering the second question before answering the first question. He uses the example of haste, stating that we must first explain what it is for something to be done hastily before we can provide an account of what haste is.

For haste and combustion the orders of explanatory priority are clear. They are clearly different. And they are clearly dictated by the metaphysical character of the thing to be explained. When explaining combustion we have no choice but to answer question 1 first. When explaining haste we have no choice but to begin with question 2. This difference in the order of explanatory priority for haste and for combustion is a reflection of the fact that they belong in different metaphysical categories. [The first] approach of answering question 1 first is appropriate only if attention belongs in a specific metaphysical category. [The second] approach of answering question 2 first is appropriate only if attention belongs in a different metaphysical category. (Mole, 2010, p.26)

Mole (2010) claims that whether we consider attention to be either a process or a phenomenon will determine the order in which these questions will need to be answered. He rejects the idea that attention is a process and, as such, he endorses the second approach. He

argues that we should answer the question ‘what is attention?’ by first providing an answer to the question ‘what is it for something to be done attentively?’. He claims that the first approach would be insufficient as attention is a phenomenon and not a process. He shows how the first approach is unable to account for phenomena by using haste as an example:

Suppose now that we are given the task of specifying the set of instances of haste. If we are allowed to use only the classificatory resources of a process-based taxonomy then we will be unable to do it. Again, the problem has nothing to do with multiple realizability. The problem is not that there are too many sorts of things that can be done hastily. The problem is that instantiating the same processes as an instance of haste is not sufficient for being an instance of haste. The set of instances of haste cuts any taxonomy by process. (Mole, 2010, p.28)

Mole (2010) claims here that specifying the instantiating processes of haste is not enough – as for example, one can run without haste. This shows that the metaphysical problem of attention is not just that there are too many processes that can be utilised for attention, but also that these processes need not always be attentional.

However, it is important to note that this is only a problem if we accept Mole’s characterisation of attention as a phenomenon as opposed to a process. If attention were a process, then it would, itself, be the process that constitutes attention.

Mole (2010) makes it clear that he is providing an adverbialist account of attention, which attempts to answer the metaphysical question by providing an answer to the question “what is it to do something attentively?”. He answers this question by stating that what it is for something to be done attentively is for it to be done with ‘cognitive unison’. In the next section, I will discuss Mole’s arguments in favour of adverbialism.

Section 6.4 – Contending Accounts of Attention

In this section, I will examine several different accounts that Mole assesses. The two main types of accounts that will be discussed in this section are: process-first accounts of attention and adverbialist accounts of attention.

Section 6.4.1 – Process-first Accounts of Attention

Mole (2010) explains that process-first accounts of attention are accounts that take the first approach, by attempting to answer the metaphysical question ‘what is attention?’ by identifying a process or set of processes. On a process-first account, attention *is* this process (or these processes). Mole (2010) makes it clear that he will be arguing against process-first accounts of attention.

He argues against process-first accounts by claiming that there are instances where processes that are associated with attention occur without attention. He supports this claim by providing evidence to suggest that the process of feature-binding occurs during unilateral neglect. He states that it has been shown that unilateral neglect patients are able to bind features within the neglected area. He cites studies by Ro and Rafal (1996) and Olk et al. (2001), which report that unilateral neglect patients show susceptibility to illusions, such as the Müller-Lyer illusion, even when relevant parts of the illusion fall within the neglected area. Mole (2010) claims that there is no attention during unilateral neglect and, since the process of feature-binding can occur during instances of inattention, he concludes that the process of feature-binding cannot be equated to attention.

However, I have the following objections to Mole’s criticisms. First, the leap that Mole has made from providing an example against one specific process (feature-binding) to using this as an example against attention as a process in general is too great. The process-first theorist may not necessarily hold that attention is feature-binding; it could be another process altogether. Evidence against feature-binding is not evidence against all process-first accounts.

Second, process-first theorists could hold that attention does not consist of a single process, but rather that it consists of a set of processes (the multiplicity view). In this case, even if feature-binding were considered to be an attentional process, it would not constitute attention by itself. Therefore, there would be no problem in Mole’s unilateral neglect counterexample, as the presence of feature-binding alone (without the other constituent processes) is not enough for there to be attention. On this kind of view, inattention in the case of unilateral neglect may actually be predicted (if the other relevant processes are not also being utilised).

Third, process-first theorists could hold that attention is multiply realisable (the multiple realisability view). On this view, attention may be constituted by different processes in different situations. Mole’s feature-binding during unilateral neglect counterexample may not be a problem for the process-first theorist who accepts this view, as feature-binding may not

be the process that constitutes attention in this situation and so, again, inattention would be the expected result. However, I do think that this view causes other challenges for the process-first theorist. They would need to provide an explanation of how and why the processes for attention are different in different situations. However, it is important to note that it would not necessarily be impossible to provide such an explanation. Furthermore, this kind of view might actually be useful for explaining the heterogeneity of attention, as it would allow for a wider range of things to instantiate attention in certain scenarios but not necessarily in others.

Finally, the process-first theorist might claim that attention is utilised in unilateral neglect. They might claim that the evidence that Mole has provided, about feature-integration occurring during unilateral neglect, along with other findings such as the success of priming effects in patients with unilateral neglect, suggests that some form of attention is being utilised.

To summarise my criticisms, Mole's criticism of the process-first theory only holds if all of the following are true. First, the process-first theorist takes feature-binding to be the key attentional process. Second, the process-first theorist does not accept the multiplicity view. Third, the process-first theorist does not accept the multiple-realizability view and fourth, the process-first theorist accepts that there is no attention during instances of unilateral neglect.

Mole (2010) acknowledges that feature-binding may not be the only process that a process-first theorist could identify attention with. He claims that he selected feature-binding as a process to examine due to the central role that it seems to play in serial search tasks, as an 'inference to the best explanation', but he acknowledges that inferences to the best explanation are not conclusive since a new explanandum can emerge. Mole (2010) leaves his criticism of the process-first theory here.

Despite his criticisms, Mole does acknowledge the positive aspects of the process-first theory.

Mole (2010) quotes Peacocke as follows:

There are [...] some explanatory consequences of the hypothesis that there is a single general kind of attention of which perceptual and sensational attention, and conscious thought and imagination are all subspecies. It is a familiar truth about attention that any one of these kinds of attention can interrupt any one of the others. Perceptual attention can be interrupted by conscious thought; conscious thought can be

interrupted by external events which capture the thinker's attention; either of these two subspecies of the occupation of attention can be interrupted by imagination; and so on. What we have here is not merely some family resemblance between varieties of conscious states, but apparently some form of competition for the exclusive use of a single limited faculty of attention. The familiar facts about attention are explained if there is a single, suitable high-level resource, drawn upon either by perception, conscious thought or imagination... (Peacocke, 1998, cited in Mole 2010, p.41)

Mole (2010) claims that it is a strength of the process-first view that it allows us to explain why different tasks may compete for our attention, and why our attention can be pulled away from one thing to another. This is because if attention is a single process, and if there is a limited capacity, then we cannot attend to everything at once.

However, it is important to note that Mole's comments here only apply to singular versions of the process-first view (which Mole calls the 'strong' version) and not to the multiplicity view. It is also important to note that process-first theory does not necessarily include capacity limitation (this is an additional element that would be required for Mole's comments to be relevant).

Section 6.4.2 – Competing Adverbialist Accounts of Attention

Mole moves on from his discussion and criticism of process-first accounts of attention in order to assess various forms of adverbialism. Mole introduces and criticises several adverbialist accounts of attention before presenting his own adverbialist account – the Cognitive Unison theory. He aims to show that the creation of a new adverbialist account is necessary as existing adverbialist accounts are insufficient.

Section 6.4.2.1 – Attitude-based Adverbialism

Mole (2010, p.7) criticises Bradley's 'attitude-based' adverbialism, in which attention results from any function that "is prompted by interest and brings about the result of our engrossment with its product". Bradley's 'attitude-based' adverbialism identifies 'interest' as the adverbial mode. Any process can be attentional so long as it is prompted by interest and results in engrossment.

Mole (2010) provides two different counterexamples in order to criticise Bradley's account. He uses these counterexamples in an attempt to show that interest is neither sufficient nor necessary for attention.

First, he attempts to provide an example of attention without interest, in order to show that interest is not necessary for attention. The counterexample that he provides is of a scenario in which there is an annoying fly buzzing around the room, irritating a person who is trying to focus on their studies. He claims that this counterexample showcases a scenario where we attend without interest, as he claims that we attend to the fly but that we are "far from interested in it" (Mole, 2010, p.47).

However, I do not think that it is obvious that we are, in fact, "far from interested" in the fly. It could be the case that we are interested in it but that we do not want to be interested in it. It is possible for people to take an involuntary interest in things that are annoying, repulsive, or distracting. Interest does not need to align with positive values. He claims that the fact that the fly is experienced as annoying proves that it is not of interest to the subject; the experience is annoying because they are not interested in the fly. However, I disagree with this claim. I think if the subject were not interested in the fly, then the experience would not be as annoying. The experience may be annoying precisely because the subject is interested in the fly, and this unwanted interest causes a clash with their wanted or intended interest (studying). Similarly to first- and second-order desires, the subject may want to focus on their positive interest (studying), but they are being distracted by a negative, unproductive, competing interest. This could explain their frustration about the situation. Therefore, I do not think it is obvious that Mole provides a counterexample that shows a case of attention without interest.

The second counterexample that Mole (2010) provides attempts to show that it is possible to have interest without attention (showing that interest is not sufficient for attention). Mole gives the example of a tired scholar who is genuinely interested in their work, but who is no longer able to concentrate due to distractors (hunger pangs, noisy children, and a radio playing downstairs).

Again, I think that whether this is a valid counterexample largely depends on whether or not we consider 'interest' to be something that is separate from 'attention'. Bradley could claim that the scholar's attention to their work is fading because their interest is being directed elsewhere (hunger, children, radio).

Additionally, Mole (2010) characterises interest as something that is fairly static. For example, by suggesting that a scholar has a set level of interest in their studies. However, interest could also be characterised as being dynamic. The level of interest that one has in their studies, work, hobbies, etc does not necessarily remain at a set level; it is something that can vary. On this kind of view, competing interests can play a role in the variance of interest levels. A person can have many interests (their wellbeing – hunger, their family – children, and their current environment/news/music – radio), but this does not mean that they can attend to all of those interests simultaneously. It is also important to remember that for Bradley, the interest needs to result in engrossment in order for it to be an instance of attention. Bradley could respond by stating that while we can have multiple interests, these interests are dynamic and they compete for our engrossment; only the interest that successfully results in our engrossment is attended to. Mole’s interpretation of Bradley is slightly uncharitable, as he does not account for Bradley’s second criterion – engrossment. Bradley could account for the tired scholar losing their interest, and therefore losing their focus on their studies, due to these competing interests.

Additionally, these counterexamples assume that we consider attention and interest to be separable in a way that Bradley does not. So, these counterexamples will not be effective against his account since he could simply state that there is interest in the fly (since there is attention to it).

Mole (2010) claims that his counterexamples show that Bradley was wrong about interest. He also claims that “any attempt to identify attentiveness with the presence of some attitude on the part of the attentive subject would be vulnerable to similar counterexamples” (Mole, 2010, p.48). He uses this reasoning to claim that personal-level attitudes in general are the wrong sort of thing to serve as an adverbialist’s attentive manner.

However, it seems that Mole is making too much of a leap by using the rejection of a specific type of attitude-based adverbialism to reject attitude-based adverbialism in general. Even if the counterexamples that Mole provided were sufficient, which I have argued in this section that they were not, they are specifically intended to separate interest from attention. It is not clear that other personal-level attitudes are separable in the same way—especially not with these specific counterexamples. Once again, Mole has been too quick to rule out an entire category based on a counterexample of one kind from this category.

Section 6.4.2.2 – Disposition-based Adverbialism and Explanation-Based Adverbialism

Mole goes on to discuss Ryle's account: disposition-based adverbialism. On Ryle's account, paying attention to something is about how the individual's activities are situated relative to their various dispositions to do that thing and to carry on doing it, or do something similar, in the face of distractors. On this account, the adverbial mode is realising the subject's disposition towards certain actions (or potentially realising their disposition in terms of future actions).

Mole also introduces White's account: explanation-based adverbialism. He claims that White is not as concerned, as Ryle is, with what a subject might have done had things been different and is instead more concerned with what a subject actually does or has done. For White, it is not about how activities relate to dispositions but rather how actual activities relate to various other things that are being done:

Simply to say that someone is attending, or paying attention, gives us no more clue as to what activities he is engaged in than simply to say that he is practicing. What 'attending' tells us is that his activities and energies, whatever they are, are directed to and fully focused on something which occupies him (White, 1964, p.7, Mole, 2010, p.48)

On White's explanation-based adverbialist account, to do something attentively is to do it in a manner that explains the outcomes (Mole, 2010). For example, one was able to successfully complete the visual search task because they were attending to the task. On this view, attention is distinguished by its explanatory power. The fact that an agent is attending can explain certain features of their performance, such as extended duration of the performance or success of the performance.

However, one criticism of White's explanation-based adverbialism is that while it may be the case that attention can explain these features, it is not the case that it always or necessarily explains them. For example, a subject can attend for brief moments or attend and still fail their task. Prolonged performance or successful performance is not necessary for attention. Additionally, a subject can perform a task for an extended period of time due to inattention (auto-pilot) or can succeed at a task without paying attention to the task. It is not clear that prolonged performance or successful performance are sufficient for attention.

Additionally, the adverbial mode on this account (doing things in a manner which explains their outcomes) is problematic. Using this type of mode does not allow us to distinguish attention from other phenomena that can also explain their outcomes. For example, to do something angrily is also to do something in a way that explains its outcomes. Doing something angrily might result in the explained outcomes of being heavy-handed when performing tasks. One might close doors too aggressively or slam cupboards shut while performing tasks. This is not an adverbial mode that is unique to attention. More concerningly, the same can be said about doing something inattentively. To do something inattentively could be to do it in a way that explains loss of accuracy or memorisation. For example, I could inattentively pour a glass of water, overfilling it. Inattention would explain the outcome in this instance. Without explaining specifically what those outcomes are and how they are unique to that specific state, it seems difficult to differentiate states (including difficulty with differentiating attention from inattention).

Section 6.4.2.3 – The No-Self-Explaining Limitation

Mole (2010, p.49) introduces the ‘no-self-explaining’ limitation, which is described as follows:

The no-self-explaining limitation: An instance of an action cannot be explained by the agent’s attention to that very action.

Mole provides the example that an instance of reading cannot be explained by the agent’s attention to that reading.

He assesses White’s and Ryle’s accounts on their ability to account for the no-self-explaining limitation. He claims that Ryle’s disposition-based adverbialism struggles to account for the no-self-explaining limitation, while White’s explanation-based adverbialism can sufficiently account for it.

White’s explanation-based adverbialism can account for the no-self-explaining limitation, as on White’s account, attending is to be doing something in a certain way (a way that can explain the outcomes) and in order to be doing something in a certain way one must already be doing that thing. So, on White’s account, the no-self-explaining limitation holds up, as on his account it is not possible to explain that one is doing an action by the fact that one is attending to that action. This is because, on White’s account, to attend is to do an action in a

certain way; therefore, the action must have already been instantiated in order for there to be attention. The instantiation of an action cannot be explained by one's attention to it (since attention occurs after the instantiation of the action).

He claims that Ryle's disposition-based adverbialist account does not adhere to the no-self-explaining limitation. This is because, for Ryle, attention is related to the dispositions of the individual. These dispositions exist before the action is carried out and therefore, on Ryle's account, attention can explain the instantiation of actions. For example, one can be reading because one has a disposition to attend to the activity of reading. Mole (2010) treats this as a problem with Ryle's account.

Section 6.4.2.4 – Summarising Mole's Assessment of Competing Adverbialist Accounts

Mole (2010) concludes by ruling out both Bradley's attitude-based adverbialism and Ryle's disposition-based adverbialism based on his criticisms, which I have discussed throughout this section.

He claims that White's explanation-based adverbialism is promising but that it might be 'too good to be true'. He expresses concerns that it might turn out to lack any real substance. One of the main reasons for this concern is that White's account is vague.

Mole (2010) states that he will use White's version of adverbialism as a rough foundation for his Cognitive Unison theory. However, he states that his account will be further detailed in order to overcome the difficulties that White's account faces.

Section 6.5 – The Cognitive Unison Theory

On Mole's Cognitive Unison theory, attention is cognitive unison. To do something attentively is to do it in a way that necessarily displays cognitive unison.

Mole explains 'cognitive unison' as follows: " α 's performance of τ displays cognitive unison if and only if the resources in τ 's background set are not occupied with activity that does not serve τ " (Mole, 2010, p.51).

Where ' α ' is an agent, ' τ ' is a task that the agent is performing, and the 'background set' is the set of cognitive resources which an agent can bring to bear in service of ' τ '.

On Mole's account, cognitive unison occurs if and only if the resources or processes that could be used in service of a task (τ) are not occupied with irrelevant activity (activity that does not serve τ).

While Mole has provided a definition for the term 'cognitive unison' here, further definitions will also be required in order to fully specify his account. He provides definitions for several relevant terms in order to ensure that his Cognitive Unison theory is specific enough to account for complex cases of attention.

The three main terms that he provides definitions for are: 'tasks', 'cognitive processes', and 'bringing to bear with understanding'. I will discuss each of these definitions in their own dedicated subsections.

Section 6.5.1 – Defining 'Tasks'

According to Mole's definition, a task is an activity that is performed with guidance from the performing agent's understanding of the activity (Mole, 2010). In order for an activity to count as a task, it must meet the following three criteria: first, it needs to be agent-involving. Second, the agent needs to have an understanding of the activity, and third, this understanding needs to guide the agent's performance of that activity.

Mole (2010) makes a clear distinction between 'activities' and 'tasks'. One can carry out an activity without meeting these criteria but not a task. Tasks are necessarily involved in attention, on Mole's account, whereas activities are not. He claims that tasks do not need to involve physical interaction with the external environment; for example, looking at a chair could count as a task as long as the agent was performing this activity in a way in which their understanding of the activity was guiding their performance.

However, I have the following criticisms of this definition: first, it creates a difficulty with accounting for attention to the things that the task concerns as opposed to attending to the task itself. For example, on Mole's account, we would not be attending to the chair itself but rather attending to the task of looking at the chair.

Second, it creates difficulties with accounting for instances of passive attention. There seem to be instances where we attend to our environment passively without utilising any form of activity guiding 'understanding' about our environment. For example, passively looking at the chair.

Finally, it creates difficulties with accounting for attention to novel tasks. There seem to be instances where we attend to novel tasks, which we do not have a pre-established understanding of (and so task-relevant understanding cannot guide our performance on these tasks).

Mole may respond to my second criticism by stating that the agent does not need to have a strong understanding, which guides performance on the task, just so long as they have some form of understanding. In instances of passive perception, the understanding could be something as basic and general as the agent having an understanding of how to view an object.

However, I would argue that this would only create further problems for Mole's account. If understanding can be as general as this, then it seems to be difficult to think of instances of activities that do not utilise any understanding at all, and therefore the majority of activities (if not all activities) could be understood as tasks. This diminishes the purpose of even creating a distinction between activities and tasks, and also diminishes the point of introducing the 'understanding' criterion to his account.

Additionally, in response to my third criticism, Mole could state that understanding is a dynamic process. He could state that even if we do not initially have an understanding of novel tasks, we come to have an understanding of them through our performance of them.

I would agree that this is possible. However, Mole would not be able to adopt this response without changing aspects of his account. As on his current account, the agent cannot be attending until that understanding is developed (as only tasks can be attended to, and understanding must be utilised in order for an activity to be a task).

While the inclusion of this understanding-guided element may lead to criticism, it is an important aspect of his account. Mole includes the understanding-guided component in order to distinguish between instances of attention and inattention. He claims instances that may seem to be attentional are not actually attentional if they are not guided by understanding. For example, seeing a chair without deploying an understanding of 'seeing', tripping over a chair, or being involuntarily reminded of a chair are all non-attentional activities according to Mole. However, seeing a chair with understanding, tripping over a chair intentionally, or voluntarily remembering a chair are all attentional tasks since they all involve the utilisation of understanding in a relevant way.

However, I would argue that these examples may instead demonstrate the importance of voluntariness, as opposed to solely demonstrating the importance of the understanding aspect.

Additionally, the understanding aspect of his account is confusing – it is not intuitively clear what the difference between ‘seeing’ and ‘seeing with understanding’ is. This is especially the case if that understanding is, or can be, as general as an understanding of how to look.

Furthermore, the exclusion of involuntary remembering (and perhaps involuntary looking/movement in the other two cases) suggests that Mole’s account may have difficulty accounting for instances of salient or passive attention. However, Mole claims that his theory can account for salience as long as our attending to these salient instances is guided by an understanding of what we are doing by attending to the salient stimuli. For example, if a sudden sound is not integrated into an understanding-guided activity, then it is merely ‘disturbing’ and not attended.

However, the difficulty here is that salience has an element of attentional capture that Mole’s account does not seem to allow for. Mole’s no-self-explaining limitation states that attention cannot provide an explanation for why the individual is performing an action and, as such, attention cannot account for instances where attention is initiated – such as in cases of salient attentional capturing.

Mole (2010) makes it clear that the unison-instantiating performance of non-task activities does not constitute attention. In order for there to be attention, there must be cognitive unison in relation to a task. He provides the example of daydreaming as one such non-task activity. Even if cognitive unison was displayed during an instance of daydreaming, that instance would not be attentional.

Section 6.5.2 – Defining ‘Cognitive Processes’

In this subsection, I will discuss Mole’s definition of ‘cognitive processes’.

Mole (2010) identifies the following key features of cognitive processes: cognitive processes are agent-level; this rules out autonomic processes such as heartbeats, etc. Cognitive processes are representation-handling; this means that cognitive processes work with and manipulate representations. Cognitive processes are things that serve (or could serve) tasks that the agent performs.

To reiterate, Mole makes it clear that attention itself is not a cognitive process, although it involves cognitive processes.

Mole explains that non-cognitive processes cannot constitute an instance of attention – even if they operate in unison. For example, a jogger may utilise a whole battery of non-cognitive processing resources in their activity of jogging, such as heart rate monitoring, breathing, blood oxygen maintenance, blood sugar maintenance, balancing, monitoring the ground for obstacles, and controlling the movement of his legs. All of these processes may be focused on the task of jogging, and so they are in unison. However, since these are not cognitive processes (according to Mole), they do not result in cognitive unison, and so they do not amount to an instance of attention.

Mole claims that: “One may jog while giving one’s attention, undividedly, to the question of what to have for dinner, and if one does pay attention to one’s jogging, it is not the state of one’s automatic processes that explains why” (Mole, 2010, p.59).

However, I would disagree with what Mole counts as ‘non-cognitive processes’. In the list given above, some of these examples seem more cognitive and attention-involving than the others. For example, one cannot stop blood sugar regulation by not paying attention to that process, but one can fail to detect obstacles on the ground by not paying attention, and perhaps one could also fail to perform specific leg movements or certain balancing techniques by not paying attention.

Additionally, it is not obvious that one really can give their complete and undivided attention to thinking about what to have for dinner while jogging. Why should we not think of this as a case of divided attention? Where a majority of attention is being paid to thinking about what to have for dinner but that some attention is still being paid to the processes involved in jogging.

Section 6.5.3 – Defining ‘Understanding’

On the Cognitive Unison account, ‘understanding’ is about having a practical sense of a task as opposed to having a full conceptual understanding of said task. Mole notes that our understanding does not necessarily involve “resources that would in fact further the performance of the task” or even the “resources that are involved in what the agent knows to be his best strategy for performing the task” (Mole, 2010, p.61).

Section 6.5.4 – Summarising the Cognitive Unison Account

After unpacking all of the terms that Mole uses for his account, it is clear that his claim, ‘attention is cognitive unison’, provides a very simplified explanation of the Cognitive Unison theory.

Attention is not just about displaying cognitive unison on Mole’s account. Mole has made it clear that there are instances of cognitive unison that are not attentional. For example, instances of unison that involve non-task activities and/or non-cognitive processes, and those that are not guided by understanding, are not attentional.

Section 6.6 - Explanatory Power of the Cognitive Unison Theory

The previous section was dedicated to outlining and assessing the core theoretical claims of the Cognitive Unison theory. In this section, I will assess the explanatory power of this theory.

Section 6.6.1 – Accounting for Task Failure

The Cognitive Unison theory can account for attention even in instances where tasks are failed. Mole accounts for the utilisation of attention in the case of two different types of task failure: ‘overlearned’ tasks and ‘impossible’ tasks. Mole uses these examples to demonstrate that a task does not need to be successfully completed in order for it to be attended to.

The first category is overlearned tasks. Mole (2010) provides the example of Bob the teenage actor, who has the task of ‘looking natural’ when he walks out onto the stage. However, Bob uses the ‘wrong resources’ and overacts and as a result fails to perform this task despite attending to the task. I will expand on this with my own example: when I was learning to kickflip, I was kicking down on the skateboard instead of out to the side, as I mistakenly thought that this was what I needed to do in order to achieve the correct rotation. I did not realise at the time that this technique is incorrect and that it almost always results in failure. Despite the fact that I failed at this task, I was still attending to it; it is just the case that I was doing so by utilising the ‘wrong resources’.

Another way that Mole (2010) accounts for failure on his Cognitive Unison theory is through attendance to impossible tasks. He gives the example of a trainee clairvoyant who is focusing

their attention on determining the correct card by using their clairvoyant abilities. Despite the fact that this is an impossible task, since there are no such things as clairvoyant abilities, the trainee clairvoyant can attend to this task.

These two examples were designed to show that success is not a necessary condition for attention and that attention can be paid to impossible tasks. Additionally, they were designed to show that the resources that are utilised do not need to be the optimal resources for the task at hand; in fact, they may even be a hindrance to the completion of the task (in the cases of Bob's poor acting skills or my kicking down instead of out).

Section 6.6.2 – Mastery and Inattentive Performance

Mole (2010) claims that his Cognitive Unison theory can also account for inattentive performance during the performance of skills that have been mastered.

He claims that mastered skills are more likely to result in inattentive performance because it is difficult to sustain attention to tasks that are well understood by the agent. The reason for this is that, on the Cognitive Unison account, a well-understood task has more resources (a larger resource pool) that need to be managed in order for cognitive unison to occur (and therefore, more resources that need to be utilised for attention). This larger resource pool makes it more difficult to achieve cognitive unison, and so experts who have achieved mastery over a certain skill are less likely to attend to their performance of that skill.

However, I am not convinced by this. It does not seem obvious that well-understood tasks actually are less likely to result in sustained attention. In fact, individuals with mastery over a certain skill may even engage in frequent periods of sustained attention. For example, athletes often have a deep understanding of their chosen sport and may often engage in long periods of sustained attention while practising or performing that sport.

Additionally, even if it were the case that mastery leads to difficulty with sustaining attention, this phenomenon could be otherwise explained. For example, this could be explained by loss of interest. Sustained attention may be strenuous to maintain for prolonged periods of time, and when an expert feels that they will no longer gain anything from attending to their performance, they may start to lose interest and decide that the effort is no longer worth it.

Section 6.6.3 – Superfluous Competence

Mole (2010) points out that accounting for superfluous competence could be difficult for his Cognitive Unison theory.

The problem is that one could use several different strategies when attending to a task. This causes an issue with how we should understand cognitive unison since there will be many different strategy-relevant resources in any given task's background set. If task X could be achieved by utilising either strategy A or strategy B, then the resources for both strategy A and strategy B are in the background set for task X. This means that for cognitive unison to be achieved the resources for both strategy A and strategy B must be in operation, even if only strategy A (or only strategy B) is being used to carry out task X. Mole (2010) provides the example of solving a mathematical problem through visualisation versus solving the same mathematical problem through calculation. In this example, all of the resources used for visualisation would need to be unified with all of the resources used for calculation, since these are both resource sets that can be used to carry out the task of solving the mathematical problem. This is obviously problematic for the Cognitive Unison account. This problem is further heightened when considering the fact that certain tasks may be performed through many varied strategies. If a task has fifty varied strategies, then fifty varied resource sets will be required to achieve cognitive unison, even if only a single strategy is actually being implemented.

However, I think that reframing tasks could help Mole to resolve this problem. Mole could resolve this by being more specific about the task, based on the strategy that is implemented. For example, allowing that using either strategy A or strategy B would result in the performance of two different tasks. On this approach, solving the mathematical problem by visualisation and solving the same mathematical problem by calculation could be considered as distinct tasks. Therefore, both resource sets would not be required for both tasks since they are distinct tasks with distinct resource sets.

Mole is firm that the 'background set' for a task should be defined as the set of resources that could be brought to bear in service of the task, and he rejects the idea that the 'background set' should be defined as the set of resources that actually are brought to bear in service of a task.

However, this does not prevent him from being able to adopt my suggested response. As dividing different strategies into different tasks also changes the set of resources that could be brought to bear in service of those, more specific, tasks.

Mole (2010) takes a similar approach to the one that I have suggested, by identifying two different types of tasks – ‘superordinate’ tasks and ‘subordinate’ tasks. Superordinate tasks are the larger, overarching, main tasks that are carried out through the completion of subordinate tasks. Subordinate tasks are the smaller tasks that are used in order to complete superordinate tasks. Mole claims that we can think of different strategies for completing a superordinate task as different, distinct, subordinate tasks.

Section 6.6.4 – Lack of Privileged Processes

Mole (2010) notes that we may intuitively consider certain cognitive processes as being more important or central to attention than others. For example, working memory is a process that seems to be central to attention. He notes that his Cognitive Unison theory does not weight cognitive processes like this. The scale of the role that each process contributes is unimportant. That is to say that all of the processes in the background set are equally weighted. What is important, on his view, is that every cognitive process in the background set is utilised in unison.

However, Mole (2010) notes that certain processes, such as working memory, can still have a special status and be viewed as important processes on his Cognitive Unison theory. This status can be demonstrated by how frequently a cognitive process is contained within a background set. A process like working memory is special/especially relevant for attention not because it does more heavy lifting or because it plays a more central role during each instance of attention, but rather because it shows up frequently across many different background sets. For example, working memory has a special status because it is required for a wide range of different tasks.

However, the idea that all processes are equally important, and that attention is not possible without even a single one of these processes, seems unintuitive. To demonstrate this, consider the following example: imagine that someone could recall vivid details of witnessing a car crash – they could tell you the license plate numbers of the cars involved, the sounds of the event, the details about the location in which it happened, etc. but they were unaware of the smells at the scene of the crash. Yet, despite having all of these experiences associated with

their witnessing of the event, they cannot be said to have attended to the event. This is because they were unaware of the smells of the scene, and since olfactory processes are something that could have been brought to bear in service of this task, the whole background set was not utilised. As such, in this example there was no cognitive unison and therefore, no attention. Without a concept of weighting, the olfactory elements are just as important for the witnessing of the car crash as the visual elements on the Cognitive Unison account.

Mole could respond to this criticism by emphasising that his account frames cognitive unison in a privative way. He could claim that cognitive unison is not about utilising the entire background set for the task at hand, but rather it is about ensuring that none of the resources in the background set are being utilised in service of another competing task. On this view, it would be possible for cognitive unison (and therefore attention) to be achieved even when some of the resources in the background set are not utilised. Mole could therefore respond to my car crash example by stating that it is fine for our olfactory resources not to be utilised, as cognitive unison is still achieved without them so long as these olfactory resources are not utilised elsewhere.

However, the problem still arises if the car crash example is adjusted so that the witness also happened to be utilising their olfactory senses on an irrelevant task (smelling perfume) while witnessing the car crash. Since olfactory resources could have been brought to bear on the task of witnessing the car crash and they were instead being utilised by another irrelevant task (smelling perfume), cognitive unison would not be achieved while witnessing the car crash. Therefore, the witness would once again not be attending to the car crash, despite having all of the other non-olfactory experiences listed in the previous version of this example.

Mole (2010) claims that there is also another way that cognitive processes can be ‘special’ (especially relevant for attention). They can be special by being “good at marshalling the support of other resources to the task that it serves” (Mole, 2010, p.68).

He provides the example of working memory as one such cognitive process. Mole (2010) states that a study by Lavie (1995) found that working memory helps subjects to remain fixed on a task (even when the task makes relatively few working memory-related demands).

However, it is not clear why this should be taken as evidence that working memory is ‘good at marshalling other resources’ as opposed to it being taken as evidence that working memory is a weighted process.

Additionally, just because a task does not explicitly demand the use of working memory does not mean that working memory is not used for that task. This is not evidence that working memory is not being used and that it is instead a ‘support to marshal other processes’. Rather, it could be the case that working memory is used in many attentional tasks regardless of whether the task explicitly demands it or not.

Furthermore, it could be the case that storage to working memory is a consequence or a result of attending as opposed to an integral part of the process of attending. This may explain why it is associated with attention so frequently and why it occurs even when the task did not demand it.

Section 6.6.5 – The Privative Character of Cognitive Unison

Mole (2010) characterises two different ways to posit the relevancy of cognitive processes: the positive view and the negative view. On the positive view, in order for something to be an instance of attention, there needs to be a presence of relevant processes, and in order for something to be an instance of inattention, there needs to be an absence of relevant processes. However, on the negative view, in order for something to be an instance of attention, there needs to be an absence of irrelevant processes, and for something to be an instance of inattention, there needs to be the presence of irrelevant processes.

He acknowledges that the current view of cognitive unison that he has set out is a ‘negative’ one. It is about the absence of irrelevant processes as opposed to the presence of relevant processes. However, he also claims that he feels that using irrelevant processes encompasses the inclusion of relevant processes. He claims that the ‘positive’ and ‘negative’ framing tend to coincide extensionally, as the only way to have an absence of something irrelevant is to have the presence of something relevant.

However, I would disagree with this point. The presence of relevance is not the only way to have the absence of irrelevance, as having nothing is a way of having an absence of irrelevance without having the presence of relevance.

Mole (2010) distinguishes between two views that are related to this positive/negative framing. The first view is the privative view, which is a negative view, in which resources do not need to be actually/currently utilised and in which cognitive unison occurs when there is an absence of irrelevant processes. The second view is the threshold view, which is a positive

view, in which resources do need to actually be utilised and in which cognitive unison occurs when there is a presence of relevant processes.

Mole (2010) favours the privative view for the Cognitive Unison theory. He provides an example, based in birdwatching, to show why he prefers the privative view. He claims that when birdwatching, if there is a period where there are no birds present, the privative view still allows for us to be attending to the task (which Mole claims is intuitive), despite the fact that no resources are being used on watching birds. Whereas, on the threshold view, there is no attention to the task if there are no birds (since resources need to actually be used in order for there to be attention).

However, I would disagree with Mole. The threshold view proponent could state that they are attending to the task of birdwatching; they are just not attending to birds. This seems intuitive as there are no birds to watch in this scenario. The problem that Mole has set out here is less about the framing that each view provides and more about how to define a ‘task’ (birdwatching versus watching birds). The threshold view proponent can still positively utilise resources to ‘birdwatch’ (task) even if no birds are present. For example, they can deploy resources to watch certain spaces or to listen for sounds, etc.

Additionally, this means that, on Mole’s account, one can be said to be attending to a task when nothing is happening. For example, if I have resources 1, 2, and 3 in my background set for jumping hurdles, then I could be said to be attending to jumping hurdles even when I am actually doing nothing, as long as processes 1, 2, and 3 are not being used for anything else. This negative framing seems to be incredibly problematic, as it allows for anything to be considered as attended to as long as the relevant resources are not being utilised elsewhere. This means that if no resources are being utilised at all, then all tasks could be said to be attended to.

Mole (2010) gives another example, demonstrating the rare occurrences where positive and negative framings come apart, in order to show why he favours the negative framing of the privative view. He provides the example of Buddhist meditation in order to demonstrate this. He claims that during Buddhist meditation, thoughts and feelings are focused on without exercising any judgement about the content of those thoughts. He claims that the negative view explains this well – by explaining this phenomenon as ‘pure cognitive unison’ free from any irrelevant processes. He claims that the threshold view cannot account for this (as there are no relevant processes being attended to) and so they must either state that there is content

or that these instances of Buddhist meditation are not attentive. Mole claims that both of these options are undesirable, and as such, this causes a problem for proponents of the threshold view.

However, I will argue that it is not clear that these meditative states lack relevant processes. Just because the thoughts that arise during these states of meditation do not have specific content does not mean that they could not be attended to in and of themselves as thoughts. A proponent of the threshold view could state that there is a positive relevant process – the sensation, which is present in the meditation process itself, of attending to these contentless thoughts. The threshold view does not necessarily rely on the thoughts having content in order for them to be attended to – as there are still relevant processes that are present. For example, the processes that are involved in the having of the thoughts themselves and the processes that are involved in the meditation process. These processes can occur regardless of whether or not the thoughts that are being attended to have content.

Section 6.6.6 – Divided Attention

Mole (2010) is clear that he wants his Cognitive Unison theory to be able to account for the following aspects of divided attention. First, he wants it to be able to account for divided attention in general. Second, he wants it to be able to account for the difficulty that is involved in the process of dividing attention, and finally, he wants it to be able to account for instances where it is impossible to divide attention.

Mole (2010) claims that, on the Cognitive Unison theory, it is possible for an agent to attend to more than one task simultaneously, without dividing their attention. He uses the previously discussed example of solving a complex math problem to demonstrate this:

Not all cases in which an agent is attending to more than one task are cases in which dividing one's attention presents its characteristic difficulties. In the mental arithmetic example that we considered [...] the agent may be said to be paying attention to the calculation of the difference between the two squares, and she may, at the very same time, be said to be paying attention to visualizing the area of overlap between them. The task of finding the difference and the task of visualizing the overlap are distinct (since an agent may be doing either one of them without doing the other). This is therefore a case in which two distinct tasks are being given attention at one time. But it is not a case where attention is *divided*, since one of these tasks is being executed as a way to accomplish the other. It is only when the agent is dividing her attention

between a pair of *competing* tasks that a distinctive difficulty arises. (Mole, 2010, p.74-75)

He claims that these two strategies really are distinct tasks because we could do them separately and that this is not an instance of divided attention, since one of these tasks is being executed in order to accomplish the other. This shows that the utilisation of subordinate tasks to complete superordinate tasks is not an instance of divided attention on the Cognitive Unison theory.

However, just because they could be carried out separately does not mean that they are separate or distinct when they are carried out together. For example, audio processes can be carried out without visual processes and vice versa, but they can also be combined into audio-visual processes. Why could it not also be the case that performing the visual strategy and the calculation strategy simultaneously leads to a new combined strategy (the visual-calculative strategy)?

Additionally, if attention is being given to one of the processes in order to “accomplish the other”, then it seems like the second process is not even being attended to, at least not as a means to resolve the mathematical problem directly. At the very least, it seems that the way in which the second process is attended to is in some way different from the way in which the first process is attended to, as the first process plays a more active role and results in the second process.

Section 6.6.6.1 – Easy Cases of Divided Attention

Mole (2010) claims that divided attention only occurs in instances where there is a pair of competing tasks.

He provides an example to demonstrate that there are instances in which it is easy to divide our attention. The example that Mole (2010) provides comes from a study by Allport, Antonis, and Reynolds (1972) in which participants were tasked with simultaneously reading piano sheet music and performing an auditory shadowing task. The results of this study showed that most participants were able to perform this task successfully without difficulty. Mole (2010) claims that his Cognitive Unison theory can account for the ease of the participants' performance on this task. He claims that ‘divided attention’ was easy in this case

because both of the tasks had distinct background sets with no overlapping processes. This meant that there was no competition between the two tasks that were used in this study.

However, it is important to note that this lack of competition also means that divided attention is not actually occurring on the Cognitive Unison theory (since competition is a necessary aspect of divided attention). Despite this, he claims that his theory accounts for this as an ‘easy’ case of divided attention. This is obviously contradictory. I would charitably interpret him as discussing this ‘easy’ case of ‘divided’ attention as a case of ‘supposed’ divided attention, where the ease of this case is actually attributed to the fact that there is no divided attention occurring at all.

Section 6.6.6.2 – Difficult Cases of Divided Attention

As previously discussed, Mole (2010) explains that some processes, such as working memory, feature frequently in a wide range of background sets. He claims that we should expect that divided attention will not be ‘easy’ across tasks that feature common processes, such as working memory, since there are much higher odds of the background sets of these tasks clashing with the background sets of other tasks.

This explanation clashes with his previous claims. Here he is suggesting that it is possible to have attention when background sets clash (which was previously the main criterion for not having attention) but that it is not ‘easy’. This is different from his previous claim that it is not possible to have attention under these circumstances.

He claims that in some cases it will be difficult to utilise divided attention and that divided attention may be ‘unstable’ in these cases.

He begins his explanation of ‘unstable’ divided attention with an example of a person riding a train, looking out of the window while listening to music through their headphones. He claims that instances like this are not stable, as if something interesting happens outside, they will stop attending to the music (and vice versa).

Mole (2010) claims that the train example is different from the piano and verbal shadowing task. On the piano and verbal shadowing task, there were distinct inputs and outputs—visual and auditory inputs and manual and verbal outputs. However, Mole claims that while the train task does have distinct inputs (auditory and visual), it differs from the previous example, as it has clashing outputs (both outputs concern occurrent thoughts and reflection).

However, it is not clear why these have to be the outputs of this task. Why could the outputs not be visual representation and visual experience in the case of looking out of the window, and auditory representation and auditory experience in the case of listening to music?

Additionally, treating general thoughts about the current environment or scenario as outputs is potentially problematic, as this could collapse the separate (non-clashing) outputs of virtually every task. For example, on Mole's piano reading and auditory shadowing task, we could say that the overarching output (or at least an output) is the subject's reflectance on their performance (thinking about how they are performing on each task). This could be considered to be a shared output for both tasks, which would result in their collapse—making it impossible to attend to them both simultaneously on the Cognitive Unison account.

Furthermore, if the goals or outputs are very broadly specified, could it not be the case that all tasks may share an output/goal? For example, the piano and auditory shadowing task, which Mole has previously classified as an instance of 'easy' divided attention, could actually be considered to be an instance of difficult divided attention by reconsidering the goals or outputs involved in these tasks. If the participants were to reflect on their performance of both of these tasks, then there would be a clash in the outputs between these tasks. Additionally, even if the participants were not to reflect on their performance of both tasks, the goal of both tasks could be described as a shared goal, since they were asked to do both tasks simultaneously as a joint task, so the goal may be a shared goal of joint performance. This clash in outputs or goals should create a difficult case of divided attention despite Mole's previous characterisation of this as an 'easy' case of divided attention and the participants' ease of performance on both tasks.

Section 6.6.6.3 – Quiescence in Divided Attention

Mole (2010) claims that it is possible for tasks with overlapping background sets to be attended to simultaneously, as long as the overlapping processes are not taken up with anything at all—as long as they are 'quiescent'. When this changes, and the processes in the overlap are not quiescent and are utilised, then it is no longer possible to share attention between these two tasks.

He claims that there are two ways for tasks with overlapping processes in their background sets to be attended to simultaneously. First, when the overlapping processes are quiescent.

Second, when the overlapping processes are active, but they serve the performance of both tasks.

However, this goes against his previous commitment that cognitive unison is about bringing to bear all of the possible resources in the background set as opposed to all of the resources that are actually utilised in the background set. Quiescent resources are still possible resources in the background set even if they are not actually utilised.

Mole (2010) claims that the reason the train rider from his previous example is able to attend to both the scenery and the music simultaneously is that they are ‘non-committal’ about attending to the scenery and/or music. He claims that this leaves some of the resources in the background set quiescent, which allows them to be free for use in other tasks.

He claims that in the case of tasks that are not non-committal (and therefore where resources are not quiescent) and that have intersecting background sets, there is no divided attention. This means that, on the Cognitive Unison theory, these instances are not genuine instances of divided attention. In these kinds of cases, attention is not paid simultaneously, but rather sequentially. Mole is satisfied with this limitation of his account. He gives examples, such as typing while correcting one’s posture and eating while having a conversation, to show that this limitation is intuitive. He claims that the writer pauses their typing, suspending their attention to it, in order to correct their posture, and that the same is true of the meal-time conversationalist.

However, it does not seem clear that these tasks have a shared background set. Thus, it is not clear why Mole could not just account for these as genuine instances of divided attention that do not clash because there is no overlap between their background sets.

Section 6.6.6.4 – Divided Attention to Composite Tasks

Mole (2010) claims that even when two tasks have essential overlapping resources, it is still possible for both of these tasks to be attended to by grouping them together into a larger task. Doing so makes it the case that the resources from both tasks are now working in unison as opposed to working in opposition.

However, if this is a possibility, then it seems that it would be possible to group all of the seemingly incompatible tasks together so that there is always unison. This seems to generally undermine Mole’s core claims about cognitive unison being activated when resources are not

being used elsewhere for another task. If a task can be defined so broadly, then this aspect of his account is devalued. Why could a single unified task not be ‘running-auditory shadowing-listening to music-visual search-reading’, and if this could be the case on Mole’s account, then why is it the case that we would likely not be able to attend to such a task? Since things have become a single task and unison can be achieved, we should have an unproblematic instance of attention on Mole’s account.

Section 6.7 – Filling in Wu’s Framework

In this section, I will briefly consider how the Cognitive Unison theory could be used to fill out Wu’s framework for a theory of attention (which was introduced in the previous chapter). The five questions outlined by this framework are as follows: the metaphysical question, the function question, the properties question, the mechanism question, and the consciousness question.

The Cognitive Unison theory provides clear answers to two of the questions in Wu’s framework: the metaphysical question and the consciousness question. However, it does not provide clear answers to the other three questions: the function question, the properties question, and the mechanism question.

The Cognitive Unison theory has provided an answer to the metaphysical question, as this is a question that Mole has directly addressed throughout his development of this account. On the Cognitive Unison theory, attention is cognitive unison. On this account, attention is characterised by a certain adverbial mode – to have attention is to do something attentively. More specifically, on this account to do something attentively is to do it in a way that necessarily displays cognitive unison.

The Cognitive Unison theory also provides at least a partial answer to the consciousness question. On the Cognitive Unison theory, attention is not necessary for consciousness, as not all instances of consciousness require attention. This aspect of the Cognitive Unison theory is demonstrated in the following quote:

If we use the word ‘attention’ in such a way that all conscious states involve attention to some degree, even those idle, tuned-out states that we naturally describe as instances of inattention, then we lose our grip on the selectivity and directedness that

our preliminary characterization of attention took to be of its essence. (Mole, 2010, p.7)

However, the Cognitive Unison theory does not provide a clear answer to the function question, as despite the fact that Mole emphasises the importance of the selectivity of attention on his account, this is not clearly mapped onto any specific function that attention might serve.

The properties of attention are also not clearly specified on this account, although it is clear that attention must involve ‘unison’ and direction towards a ‘task’.

Finally, the Cognitive Unison account is silent about the specific mechanisms that are involved in attention. On this account, it could be argued that there is some kind of broad mechanism that allows for the creation of unison; however, such a broad mechanism is not identified. It also seems that there may be mechanisms involved in the detection of such states of unison and of relevancy (and/or irrelevancy) of certain resources to the task at hand. However, the account does not identify any specific mechanisms that are necessary for attention.

Section 6.8 – Compatibility with Monotropism

In this section, I will consider whether the Cognitive Unison theory of attention would provide a suitable foundation for a revised account of Monotropism. I will do this by comparing the compatibility between the core claims of the Cognitive Unison theory and the core claims of the extant account of Monotropism. I will argue that there are two main points of incompatibility between the extant account of Monotropism and the Cognitive Unison theory.

The first point of incompatibility is that the answer that the Cognitive Unison theory provides to the consciousness question is at odds with the characterisation of the relationship between attention and consciousness on the extant account of Monotropism. Although this relationship is never explicitly specified by the extant Monotropism account, it is suggested on this account that attention is necessary for consciousness – this allows for a straightforward interpretation of inattentional blindness in autism. However, as I have discussed in the previous section, on the Cognitive Unison theory attention is not necessary for consciousness.

The second point of incompatibility is a more serious issue. Adopting the Cognitive Unison theory as the foundation for a revised account of Monotropism would significantly blur the lines between autistic and non-autistic attention styles. This would seriously diminish the explanatory power of such a revised account of Monotropism. This is due to the fact that the way the Cognitive Unison theory proposes attention to be in general is too close to the way that the extant account of Monotropism proposes attention to be in autism. As I have already discussed in this chapter, the Cognitive Unison theory struggles to account for instances of divided attention, and Mole claims that this is unproblematic, as many cases which seem to be divided attention are really just instances of sequentially distributed attention. This characterisation of attention, which was intended to explain attention in general, is too close to that of the monotropic processing style that was intended to explain attention in autism on the extant Monotropism account. In order for Monotropism to have explanatory power, there needs to be a clear distinction between the attentional style within autism and the attentional style outwith autism. Taking the Cognitive Unison theory as a foundation reduces the clarity of that distinction.

It is for these reasons that I will not be adopting the Cognitive Unison theory of attention as a foundation for a revised account of Monotropism.

Section 6.9 – Chapter Conclusion

In conclusion, in this chapter I have introduced and discussed Mole's Cognitive Unison theory of attention. The core claim of this theory is that attention is cognitive unison. This account takes an adverbialist approach to the metaphysical question of attention, answering the question 'what is attention?' by considering what it is to do something attentively. I have explained that Mole has argued that to do something attentively is to do so with cognitive unison.

I have ultimately argued that the Cognitive Unison theory provides an unsuitable foundation for a revised account of Monotropism for several reasons. First, because of the general problems that this account faces, which I have outlined throughout this chapter. Second, because aspects of this account are incompatible with aspects of the extant Monotropism account. The first point of incompatibility is that attention is not necessary for consciousness on the Cognitive Unison theory; this clashes with the straightforward explanation of inattentional blindness on the extant Monotropism account. However, a more important point

of incompatibility is that the characterisation of attention in general on the Cognitive Unison theory is too close to the characterisation of attention in autism on the extant Monotropism account. Taking the Cognitive Unison theory as a foundation for a revised account of Monotropism would diminish its explanatory power by making it more difficult to account for the differences in attentional styles between autistic and non-autistic individuals.

Chapter 7 – The Priority Structure Account of Attention

Section 7.1 – Chapter Introduction

The previous two chapters have been dedicated to the assessment of two philosophical theories of attention: the Selection for Action theory (chapter 5) and the Cognitive Unison theory (chapter 6). In this chapter, I will introduce the final philosophical theory of attention that I will be discussing in this thesis: the Priority Structure account. The primary texts that I will be discussing in this chapter are Watzl's 2017 book 'Structuring Mind: The Nature of Attention and How it Shapes Consciousness' and his 2023 paper 'What attention is. The priority structure account'. I will first discuss Watzl's arguments against reductionism (in section 7.2) before going on to outline the Priority Structure account (in sections 7.3 – 7.6) and provide an assessment of it (in section 7.7). In section 7.8, I will defend the Priority Structure account from criticisms that have been made by Wu. Finally, in section 7.9, I will briefly consider how the Priority Structure account could fill in the theoretical framework proposed by Wu (which I introduced in chapter 5). I will ultimately conclude that the Priority Structure account provides the best foundation for a revised account of Monotropism.

Section 7.2 – Arguments Against Reductionism

Watzl (2017) claims that the problem of attention (which I have outlined previously in chapter 4) arises from the fact that the majority of research on attention does not address the metaphysical question 'what is attention?'. He claims that most research in this area is instead focused on the effects of attention and the mechanisms that are associated with attention.

I will argue that despite the wealth of scientific knowledge about the behavioral and cognitive *effects* of attention and about its neuronal and computational *underpinnings* contemporary science does not provide a satisfactory account of the *nature* of attention. We lack a scientific account of what attention *is* even though we know a lot about what attention *does* and what *enables* it. (Watzl, 2017, p.14)

Watzl (2017) claims that our examination of the associated effects and mechanisms of attention alone has not (to date) allowed us to reach a clear consensus about what attention is

– even at a basic level that could serve as a starting point for answering the metaphysical question.

Attention is described as anything from a capacity limitation, to a quantity-like resource, a specific neuronal or cognitive process, and a manner in which other neuronal and cognitive processes are going on. It is hard to even get a grip on what general kind of thing attention is (or maybe not a “thing” but a property, limit, or manner of being). (Watzl, 2017, p.21)

He emphasises the fact that the problem of attention has serious potential consequences for empirical research efforts on attention. Without a clearer understanding of what attention is, it is not clear that empirical tests that are designed to study attention really are studying attention. He claims that in order to combat this issue we need to have a better understanding of the nature of attention – we need to provide an answer to the metaphysical question. He emphasises this as follows:

The central problem in the study of attention is not that we are lacking data, that we need better brain scans, more neurophysiology, or better computational models. The central problem is: what is the study of attention actually studying? What is the nature of attention, its object of study? (Watzl, 2017, p.22)

Watzl (2017) argues that we should not use reductionism as a foundation for answering the metaphysical question. He claims that attention cannot be reduced down to a single neuronal mechanism. He emphasises that different forms of attention have different neuronal or computational mechanisms – there is no single neuronal mechanism associated with attention in general. He demonstrates this by listing different types of attention (shown in the list below) and suggests that it is unlikely that all of these types of attention share a single mechanism:

- Overt / covert attention – Involvement of bodily movement / no involvement of bodily movement.
- Focal / global – Narrow / diffuse

- On-off / graded – binary / non-binary
- Endogenous / exogenous – internally controlled / externally controlled
- Voluntary / involuntary – controlled by the subject's intentions / not controlled by the subject's intentions
- Perceptual / executive – sensory systems / central processing systems
- Disengaging / shifting / engaging – withdrawing attention / switching attention from one stimulus to another / focusing on a new stimulus.

Watzl (2017) continues his argument against reductionism by setting up a two-horned dilemma. He claims that the reductionist has two main options for providing an answer to the metaphysical question. The first option is to equate attention to a specific mechanism, and the second option is to equate attention to a general, broader mechanism. He explains that choosing either option is problematic for the following reasons.

First, if the reductionist attempts to provide a specific mechanism for attention, then, Watzl (2017) claims, they will run into the problem that there is strong empirical evidence to suggest that specific mechanisms that are associated with attention are context-relevant. There are cases where these mechanisms operate during instances of attention and other cases where these same mechanisms operate without attention. Additionally, different mechanisms seem to be associated with different forms of attention. This makes it extremely difficult for the reductionist to provide a specific mechanism for attention that can sufficiently answer the metaphysical question. However, if the reductionist instead chooses to provide a general mechanism for attention, then he claims that they will be left with the problem that there is currently no evidence for such a general mechanism of attention. Watzl (2017) claims that it is highly unlikely that there is such a general mechanism of attention. He claims that the only thing that unites these various mechanisms is that they are all associated with attention—there is no other unifying reason or broad mechanism present across all instances.

In the next section, I will introduce Watzl's own theory, which provides an answer to the metaphysical question – the Priority Structure account of attention.

Section 7.3 – The Priority Structure Account of Attention

Watzl answers the metaphysical question by providing an explanation of the nature of attention. He describes attention as an overarching process that organises other mental processes. Attention is not a mental process, like perception, but rather it is more like a ‘meta-process’ that structures these mental processes. The role of attention on the Priority Structure account is to structure these other processes in relation to the subject’s priorities. More specifically, attention is the activity of creating and regulating priority structures (Watzl, 2017).

Section 7.3.1 – The Features of Attention

Watzl (2023) claims that attention must have the following four features: it must concern ‘activated’ aspects of the mind, it must involve ‘guidance’, it is ‘meta-representational’, and it orders things in a way that is ‘usable’.

Feature 1 – ‘activated’

Watzl (2023) claims that attention only concerns what is currently active; this is a key aspect of attention that separates it from other mental processes or mental states such as memory, desire, and belief. He emphasises that attention is an ‘occurrent’ aspect of the mind.

Occurrent states are states that are currently active. An example of an occurrent aspect of the mind might be thinking about a friend at this exact moment in time, while an example of a non-occurrent aspect might be knowing that friend’s name even when you are not actively thinking about them.

Feature 2 – ‘guided’

Watzl (2023) claims that attention is a process that can be influenced and guided by a subject’s motivational systems. He specifies the following relevant guiding states that can guide attention in humans:

- General world knowledge and meaning
- Intentions, goals and plans

- Selection history
- Prior reward and associated value
- Emotional factors
- Executive control

Feature 3 – ‘meta-representational’

Watzl (2023) claims that attention is not a representational state nor a motivational state. It is because of this that attention cannot be subject to accuracy or satisfaction fulfilment criteria – it cannot be accurate or inaccurate, nor can it be satisfied or unsatisfied. He claims that the role of attention is to organise information, and that while attention may organise states that have representational content by highlighting or prioritising them, attention itself is not representational.

Feature 4 – ‘usable’

Watzl (2023) claims that attention organises the mind for use in action, cognition and memory formation. He claims that attention is used in action selection, belief formation, and in the formation of some types of conscious experience (this will be further discussed in section 7.6 of this chapter).

Section 7.4 – Attention is an Activity

Watzl (2017) further specifies the nature of attention by specifying the type of process that attention is – on the Priority Structure account, attention is an activity. He claims that in order for a process to be an activity, it must have a specific temporal shape and it must also include subject-level guidance. He states that attention is an activity, as it meets these criteria. He claims that attention is an unfolding process with the necessary temporal form for an activity (as opposed to a mental state or a single mental event) that is influenced by subject-level guidance. He claims that attention must be understood as an activity in order to account for the often voluntary nature of attention.

An opponent of the Priority Structure account could argue that describing attention as an activity would not account for passive attention where subject-level guidance (such as initiation, control, or monitoring) is not present. However, it is important to note that subject-level guidance is not the same as conscious or voluntary control on the Priority Structure account. On this account, an instance of attention can be guided by subject-level guidance without conscious or voluntary control from the subject. Examples of subject-level guidance (from the list provided in the previous section) that can guide attention passively, without conscious or voluntary control, include general world knowledge, goals and plans, selection history, and associated value. For example, if I am camping in the woods and I know that I will need to make a campfire later on, then these forms of subject-level guidance can make firewood more salient to me—even if I am not actively engaged in the task of looking for firewood at that moment in time. What appears to be salient to a subject can be influenced by guidance even if the instance of attention is not brought about or maintained intentionally. In this way, passive attention could be considered to be an activity according to Watzl since even passive forms of attention can involve subject-level guidance.

Watzl (2017) claims that “... attending remains a subject’s activity even when it is not controlled by the subject and not guided by her executive control network...” (Watzl, 2017, p.69).

In fact, Watzl (2017) claims that attention will be controlled by psychological salience unless the subject utilises executive control in order to override the control of psychological salience. This shows that the Priority Structure account not only allows for psychological salience but it also credits psychological salience with a major role in determining the default targets of attention.

Additionally, Watzl (2017) claims that all activities are guided by subject-level states. Since he considers even unintentional instances of attention as activities, he considers there to be subject-level guiding states even in unintentional instances of attention. He claims that the guiding states for instances of unintentional attention are ‘primitive guiding states’, as opposed to complex guiding states (such as goals or intentions). This allows Watzl to account for instances of attention in creatures without executive control systems. These creatures may utilise unintentional attention even if they cannot utilise more complex forms of attention such as intentional attention.

Furthermore, Watzl (2017) claims that attention can fluctuate between being voluntary and involuntary. He argues that it is because of this that attention is best understood as a mental process (as opposed to a single mental event or a mental state), as this allows for an explanation of the fact that the voluntary status of attention is often dynamic. He claims that considering attention as a single mental event would not capture this aspect of attention, as events are static – they are either voluntary or involuntary. Watzl (2017) also claims that this aspect of dynamic voluntariness is not accounted for by considering attention as a mental state. This is because mental states cannot be performed voluntarily or involuntarily, since they cannot be performed at all.

Section 7.5 – Priority Structures

So far, I have explained that the Priority Structure account answers the metaphysical question: ‘what is attention?’ by specifying that attention is the activity of creating and regulating priority structures. However, in order to understand this answer, an explanation of priority structures must be provided. What does it mean for attention to be the activity of creating, maintaining, and changing priority structures? What are priority structures? In the previous section, I have outlined Watzl’s arguments for considering attention as an activity. In this section, I will discuss the nature of priority structures.

While Watzl (2017) has argued that attention is not itself a mental state but an activity, he argues that attention is partially constituted by mental states that he calls priority structures. He explains priority structures with the following analogy:

Think of your current mental life as being like a newspaper. Just as the newspaper contains various stories and reports, your mental life contains various mental states, events, and processes. You see certain things, hear certain things, have certain feelings, thoughts, and emotions. Attention, on the view being developed, is not another such element in the paper. It concerns the placement of the stories. Is something a headline or a frontpage article? Or is it a story in the fine print, pushed to the back pages? Just like the newspaper is structured into front and back, so your mind is structured into what is of top priority and what is deprioritized. Priority structure is an organization of the mind. It contains states, events and processes, perceptions, thoughts, and emotions as parts. Priority structure is about how a mental life is formed from its parts. (Watzl, 2017, p.70)

Watzl (2017) specifies that priority structures are not the results of attending to something. It is not the case that attending is a process that is separable from priority structures that merely results in specific priority structures. Rather, it is the case that priority structures are a constituent part of what it *is* to attend to something on the Priority Structure account. Watzl compares this to the fact that a leg movement is a constituent part of the process of walking. He claims that “You can’t have priorities, unless you are prioritizing. And you can’t prioritize, unless you have priorities. State and activity are interdependent.” (Watzl, 2017, p.70).

Watzl (2017) claims that priority structures consist of both structural elements and a structural relation.

Section 7.5.1 – The Structural Elements of Priority Structures

He claims that the elements of priority structures are psychological elements as opposed to worldly elements. For example, in paying attention to a red spot on the wall, what occupies the attention of the subject is the ‘seeing of the red spot’. This is because attention consists of regulating priority structures – what occupies attention are the psychological elements, such as mental states, that are contained within the priority structure. However, one can attend to worldly elements by occupying their attention with the related psychological elements. For example, one attends to the red spot when the ‘seeing of the red spot’ occupies one’s attention.

One benefit of taking this approach, which Watzl (2017) highlights, is that it allows the Priority Structure account to account for instances of internally and externally directed attention—even in cases where internally directed attention has no clear external object, such as in the case of headaches or thinking about projects.

Watzl (2017) explains that priority structures can consist of a wide range of mental states, events, and processes. However, he claims that there are limitations to what can be included within priority structures; he specifies these limitations as follows. First, only subject-level mental states, mental events, or mental processes can be included in priority structures. Non-subject-level processes, such as autonomic processes (for example, heart rate), are not included in priority structures. However, thoughts about these processes or actions related to them (such as thinking about or measuring one’s heart rate) can be included in priority structures. Second, only occurrent mental states, mental events, or mental processes can be included in priority structures. This is related to the ‘active’ feature of attention, which I have

previously outlined in this chapter. This helps to distinguish attention from other mental processes such as memory.

It is also important to note that because other mental states, events, and processes are the constituent elements of attention, attention is dependent upon these mental states. On this account, it is not possible to have attention without them.

Section 7.5.2 – The Structural Relations of Priority Structures

Watzl (2017) then discusses the structural relations between contents within priority structures. He rejects the notion of relationships that have quantifiable differences in value. For example, he rejects the notion that ‘priority A’ in the priority structure has a value of ‘4’ while ‘priority B’ in the priority structure has a value of ‘8’. He also rejects the notion of relational quantifiable values, for example, that ‘priority A’ in the priority structure has half as much value as ‘priority B’ in the priority structure.

The main reason that Watzl rejects these quantifiable relations is to account for instances of intransitive priority orderings. These are instances where a subject has a seemingly paradoxical order of priorities – this is based on the notion of intransitive preferences. In the case of intransitive preferences, a subject may state that they prefer A over B and B over C but that they prefer C over A. For example, they prefer cats over dogs and they prefer dogs over lizards but prefer lizards over cats. The fact that they prefer cats over dogs and that they prefer dogs over lizards has secured cats at the top spot in first place, while dogs are in second place and lizards are last in third place. The subject then breaks this sequence by saying that they prefer the thing that is in last place (lizards) over the thing that is in first place (cats).

Watzl claims that assigning quantifiable values makes cases of intransitive priority ordering impossible. For example, if cats have a value of (10) in first place and dogs have a value of (8) in second place and lizards have a value of (5) in third place, then lizards cannot be preferred to cats if this value is what represents that preference (since 5 is a number of lesser value than 10).

Watzl claims that in order to allow for cases of intransitive priority ordering, the Priority Structure account needs to adopt a different notion of priority relation—weak priority ordering. Weak priority ordering is based on the notion of weak preference from decision

theory and economics. Watzl provides the following example of a weak preference: “a subject *weakly* prefers ice cream to broccoli, if she likes ice cream at least as much as broccoli (maybe she likes it more, maybe just the same).” (Watzl, 2017, p.76).

This is a comparative, non-quantifiable, relational structure – wherein the elements of a priority structure are compared directly against each other in order to determine their level of priority. It is because of the fact that this is a non-quantifiable relational structure that it can account for cases of intransitive priority orderings.

To summarise, priority structures are composed of two main parts: the elements of the structure and the relational structure between those elements. I have explained that, on the Priority Structure account, the elements of the priority structure are psychological elements (mental states, events, and processes). I have also explained that Watzl rejects quantifiable relational structures in favour of weak priority relational structures in order to account for instances of intransitive priorities.

Section 7.6 – Attention and Consciousness

In this section, I will outline the relationship between attention and consciousness on the Priority Structure account of attention.

On the Priority Structure account, attention is necessary for consciousness; without attention, there can be no conscious experience. This is due to the fact that, on this account, priorities are necessary for subjective perspectives, which are in turn necessary for conscious experiences. This is demonstrated in the following quote by Watzl (2017):

conscious minds without attentional organization are *impossible*. Without priorities, subjects wouldn't have a subjective perspective on the world, and without having a subjective perspective on the world they wouldn't have conscious experience. So, no subjects (at least: no conscious subjects) without priorities. (Watzl, 2017, p. 108)

However, it is important to note that, while attention is necessary for consciousness on the Priority Structure account, Watzl does not think that attention should be equated with consciousness. He does not think that the metaphysical question of attention should be answered solely by considering the relationship between attention and consciousness. He argues against views that take this approach, such as the appearance view of attention. The

appearance view is a broad view that claims that attention has some specific effect on our conscious experience (on appearances) and that this provides the answer to the metaphysical question. The appearance view can take many different forms, depending on which specific effect on conscious experience is specified. For example, a popular form of the appearance view is one that specifies ‘sharpness’ as the appearance effect that attention has on conscious experience. On this view, attending makes conscious experience sharper. For example, to visually attend to an object is to make its features sharper or perhaps to increase its resolution. This provides the following answer to the metaphysical question: ‘attention is what makes conscious experience sharper’. This approach directly equates attention with some specific effect that attention has on conscious experience.

Watzl (2017) makes three main arguments against the appearance view. First, he argues that it is difficult to identify a specific appearance effect of attention. Second, he argues that the appearance effects that attention has on consciousness can be replicated without attention. Third, he argues that the appearance view does not explain the full range of effects that attention has on consciousness. This is to say that attention affects consciousness in other ways, beyond just altering appearances. I will further explain these three arguments below.

Watzl’s (2017) first argument is that it is difficult to identify a specific appearance effect that attention has on conscious experience. He claims that this difficulty is due to the heterogeneity of attention. He demonstrates that there is heterogeneity at the level of general appearances and also at the level of specific appearances. First, he argues that it is difficult to identify a broad appearance effect of attention by stating that it is possible to attend to a wide range of object and feature types across a wide range of different modalities. For example, there are many features such as colour, movement, and timbre that could be the appearance effects of attention. There are also many different modalities within which attention operates, such as vision, audition, or touch (and furthermore there is even non-perceptual attention in cognition, etc.). Watzl (2017) claims that it would be difficult for proponents of the appearance view to identify a single appearance effect that is shared across all (and only) instances of attention. Second, he argues that it is also difficult to identify a specific appearance effect of attention. He claims that even if we were to only consider intramodal instances of attention (for example, to only try to identify an appearance effect for visual attention and perhaps to identify a different appearance effect for other modalities), that this would also be difficult. This is because there is still widespread heterogeneity even at the intramodal level. There are many forms of attention even within visual attention, and there

are many different appearance effects that relate to these forms of attention. For example, size, colour, saturation, resolution, etc. He argues again that none of these effects seem to be shared across all (and only) instances of attention—even when the options are narrowed (for example, by specifying visual attention only). He concludes that when attempting to specify an appearance effect of attention, “We find diversity instead of unity.” (Watzl, 2017, p.172).

Watzl’s (2017) second argument against the appearance view is that any appearance effects that proponents of the appearance view may attempt to equate attention to can be replicated without the use of attention. This shows that attention is not necessary for these appearance effects and that these appearance effects alone cannot sufficiently constitute attention.

Appearance effects such as brightness, sharpness, or loudness can be altered by increasing/decreasing the related values in the external stimulus without the subject changing their attention to the stimulus. For example, an experimenter can externally increase the values of a stimulus related to its perceived brightness in order to make it appear brighter to the subject without the subject changing their attention to the object. However, it is important to note that it is possible for attention to have appearance effects on the Priority Structure account. Watzl (2017) criticises the appearance view for using these appearance properties to answer the metaphysical question (and thereby stating that attention always has some specific appearance property), not for suggesting that it is possible for attention to have appearance effects.

Watzl’s (2017) third argument against the appearance view is that it does not account for the full range of effects that attention has on consciousness. He argues that attention affects consciousness in more ways than just creating differences in phenomenal qualities; he claims that attention also affects phenomenal structure. This is something that the appearance view does not account for. He argues that there is a phenomenal structure of consciousness that is affected by attention. On this view, attention is dynamic, and as such the elements that are included, their relationship, and their order of priority are in flux; the experience of the phenomenal structure may also include the experience of this fluctuation. The appearance view fails to account for the phenomenal effects that attention has on the structure of consciousness.

I will argue that the appearance view fails. But it fails not because there is some experiential episode, some part of consciousness, like a headache or nausea, that it does not cover. What it does not cover is how the parts are put together. Whether, for

example, the apparent sound of the saxophone is at the center or at the periphery of your conscious perspective. It is true, as many proponents of the appearance view have suggested, that when we aim to focus attention on our experience of a sound or the sensation of red, we focus attention on the apparent sound or the redness we are aware of. But exactly when we do that, there is also the shape of our conscious perspective created by putting that apparent sound or apparent redness at the center of our attention. And *that*, I will argue, is not captured by the appearance view. (Watzl, 2017, p.163)

Watzl (2017) claims that the Priority Structure account does not face these same problems when attempting to account for the phenomenology of attention. This is because appearance effects are only correlates of attention on the Priority Structure account, while priority structures are constituent parts of attention. “While priority structures may correlate with certain appearances, they are not constituted by such appearances. A certain distribution of attention therefore never *entails*—as part of what it is—any specific appearances.” (Watzl, 2017, p.181).

Additionally, the Priority Structure account can also account for the structural phenomenology of attention. This phenomenology is associated with priority structures and cannot be replicated without them. The structural phenomenology of attention cleanly maps onto the structural relations of the priority structures. What it is to have an experience of structural phenomenology just is to have a priority structure that reflects that structure. For example, to have a structural experience where 1 more than 2 more than 3 more than 4 just is to have a priority structure where 1 more than 2 more than 3 more than 4. While having an experience of some feature-based appearance quality, such as ‘redness’, could be caused by multiple attentional or perceptual states, having a specific structural experience such as red more than blue more than green more than yellow just is to have a priority structure where red more than blue more than green more than yellow. On this view, there can be many different phenomenological profiles even for instances of attention that concern the same range of objects or properties. For example, ‘red more than blue more than yellow’ would have a different phenomenology from ‘yellow more than red more than blue’—because the phenomenal structure is taken into account.

The reason we could find a replica for each effect of attention on the appearances was that all of these *are* mere effects or correlates of a certain distribution of attention.

Therefore, it was possible to replicate the relevant effects in a scenario with a different distribution of attention. Priority structures, by contrast, are not mere effects or correlates of attention. They are partially constitutive of attention. A phenomenal contribution of a certain distribution of attention that is constituted by priority structures cannot be replicated with a different distribution of attention. (Watzl, 2017, p. 186)

Watzl (2017) continues on to further specify the relationship between attention and consciousness on the Priority Structure account. He claims that, on this account, attention is necessary for accessible phenomenal consciousness but that it is not necessary for other types of consciousness. This means that attention is necessary in order for one to have awareness of their conscious states in a way that would allow them to utilise information from these states for action and reasoning ('Access Consciousness', 2013).

He also claims that attention is not sufficient for consciousness on the Priority Structure account, as while attention can shape conscious experience, it is possible for there to be instances of attention without consciousness. This is even the case for top priorities within a priority structure, as otherwise this would imply that there is some sort of 'border' or 'threshold' relating to the degree of priority at which consciousness comes about.

If *some* or *quite a bit* of attention (a fairly high position in a priority system) is not enough for consciousness, then it would seem like a bit of magic why rising *all* the way to the top would suddenly necessitate that the lights are on. (Watzl, 2017, p.246)

Watzl (2017) presents two main points of evidence for his claim that attention is not sufficient for consciousness. First, that there are cases of attention in blindsight. Second, that Norman, Heywood and Kentridge (2013) conducted an experiment that demonstrated the same object advantage in cases where the subject did not have conscious awareness of the object. These examples are intended to show that there are instances of attention without consciousness.

Another potential example of consciousness without attention comes from James, who claims that distraction is the opposite of attention (Watzl, 2017). Watzl notes that since there is something it is like to be distracted, distraction could be considered to be a case of consciousness without attention. However, Watzl claims that James has incorrectly defined distraction as the opposite of attention. Watzl instead claims that distraction is a form of

attention – a specific way of distributing attention. It is a state in which attention is diffuse, where everything is of equal priority (or near enough)—it is the opposite extreme of having a spiky profile of attention.

Watzl (2017) also rejects the focal necessity claim, which is the claim that the subject can only be conscious of what is at the focus of their attention. He claims that the fact that about 30% to 60% of participants in inattention blindness tasks report seeing the unexpected stimulus shows that we can be conscious of things that are outwith the focus of our attention. He claims that “The inattention blindness results show that we need focal attention for consciousness more often than we thought; they clearly do not show that such focal attention is necessary.” (Watzl, 2017, p.250).

He instead defends a view that he calls the ‘Priority System Necessity’ view, which he outlines as follows:

A mental state, process, or activity is phenomenally conscious, only if it is a psychological part of (one of) a subject’s priority system(s). [...] According to the priority system necessity view, phenomenally conscious states must be caught by the net of priority relations that connect some of the subject’s mental states. A state that is fully outside any of the subject’s priority systems will not be phenomenally conscious. (Watzl, 2017, p.251)

Watzl (2017) also argues against the independence claim, which is the claim that attention and consciousness are two distinct, independent, brain processes. He argues that there is a strong correlation between the level of attentional priority that something receives and the degree to which that thing is phenomenally accessible to the subject.

There is a strong correlation between attentional priority and *accessible conscious awareness*. If a mental state occupies more of the subject’s attention, she is at least more likely to be able to *access* its phenomenology. This correlation is hard to explain on the independence view. If attention and consciousness were underpinned by two entirely distinct brain processes, then we would not expect a correlation even between the accessibility of consciousness and attention. (Watzl, 2017, p.252)

Watzl (2017) argues that this is something that affects gist perception, where the level of attention given affects the likelihood that the subject will have access to a phenomenal experience of the gist. Paying less attention makes it less likely that the subject will have an accessible phenomenal experience of the gist. He also claims that this is even the case for stimuli that are typically considered to be naturally salient, such as faces, and that this can also affect the pop-out effect. He demonstrates this by claiming that hemilateral neglect patients are more likely to neglect stimuli that are closer to the midline of the ‘bad’ hemifield and more likely to prioritise stimuli that are as far as possible into the ‘good’ hemifield. He claims that “There is a general correlation between attentional priority and accessible phenomenal consciousness. The less prioritized, the less likely that the subject is able to notice being in a phenomenal state.” (Watzl, 2017, p.257)

The reason why attention is necessary for consciousness on the Priority Structure account of attention is that consciousness requires a subjective perspective, and attention is what structures the subject’s perspective (Watzl, 2017).

The argument starts with the idea that a subject’s conscious experience consists in her unified and subjective point of view or perspective. Attention shows, it is argued, why consciousness is essentially unified, perspectival, and subjective. On the resulting view it is essential to consciousness that a subject takes a stance in her own conscious experience. Consciousness consists of the internal unity of a qualitative and subjective character. The qualitative character of consciousness pertains to the way the world appears to the subject. The subjective character of consciousness pertains to a subject’s actively structured point of view on those appearances. (Watzl, 2017, p.259)

On this view, consciousness is a dynamic process that necessarily involves the perspective of the subject. The subject can take a stance in their conscious experience by pushing some qualities into the centre or the periphery, and by taking such a stance they form their point of view within that conscious experience.

In conscious experience, subjects *are actively taking a perspective* on an apparent world. Consciousness has a center-periphery structure that cannot be reduced to the structure of appearances. Our structured experience is an ongoing process and not an atemporal state: guided through phenomenal saliences one moment flows into the next. (Watzl, 2017, p.259)

In this section, I have outlined the Priority Structure account of attention. In the next section, I will assess the strengths and weaknesses of this account.

Section 7.7 – Assessment of the Priority Structure Account

Section 7.7.1 – Strengths of the Priority Structure Account

In this section, I will discuss the following strengths of the Priority Structure account: overcoming the heterogeneity problem of attention and replacing capacity limitation.

Section 7.7.1.1 – Overcoming the Heterogeneity Problem of Attention

The most significant strength of the Priority Structure account is that it addresses the heterogeneity problem of attention. It is able to address the heterogeneity problem by providing an account of attention that is broad enough to account for the many varied forms of attention, while also providing an explanatory picture that shows how these forms of attention are united. On this view, what all forms of attention have in common is that they all involve the organisation of mental states within priority structures.

Additionally, the Priority Structure account is also able to provide an explanation of why attention is so heterogeneous; this is to say that it explains why there are so many varied types of attention. On this account, attention is heterogeneous because there are many different ways that priority structures can be formed or instantiated, which in turn leads to many different types of attention.

Another, related, benefit of the Priority Structure account is that it can account for attention within a wide range of modalities. The Priority Structure account can account for an instance of attention as long as there are appropriate corresponding mental states that are being organised within the priority structure. For example, it can account for perceptual attention (across all sensory modalities), cognitive attention, emotional attention, and attention related to action. The Priority Structure account of attention provides a broad explanation that unifies these types of attention—it does not need to identify attention with any single one of these types. It can account for the fact that attention is sometimes perceptual and sometimes

cognitive. It can account for the fact that attention is sometimes related to action, and sometimes it is not.

The fact that the Priority Structure account addresses the heterogeneity problem means that it also has the potential to unite a wide range of mechanisms and neural correlates that are associated with attention. Difficulties with identifying a single mechanism of attention are unproblematic for this account, as on this account the identification of such a mechanism should not be expected. This is because there can be many different types of attention (depending on the contents of the priority structure and the relationship between these contents), which could all have distinct neural correlates and measurable effects.

Section 7.7.1.2 – Replacing Capacity Limitation

Another strength of the Priority Structure account is that it can provide an alternative explanation that can account for the core claims of the capacity limitation theory without being committed to the problematic aspects of this theory.

The Priority Structure account accounts for attentional limitations through the notion of prioritisation. There is a trade-off in the process of prioritising: the more things that are simultaneously prioritised, the less priority each thing can receive. Watzl claims that attentional limitations are to be expected, as prioritising is a process of selecting and organising some things over others, leading to some things receiving more attention than others. Prioritisation requires a limit, or else it risks becoming defunct. This shows that the Priority Structure account is able to provide an explanation of attentional limitation without needing to rely on the problematic notion of attentional resources (as the capacity limitation account does).

Additionally, the Priority Structure account is able to provide a broader account of the function of attention. On the Priority Structure account, the function of attention is not solely to provide a limitation or to gatekeep for consciousness (as is the case for the capacity limitation account). Attention also functions to organise the mind more broadly, by organising a wide range of mental states, events, and processes within priority structures.

Section 7.7.2 – Weaknesses of the Priority Structure Account

As I have previously discussed in section 7.5.2 of this chapter, Watzl (2017) rejects the notion of using quantifiable values for priority ordering in order to allow for cases of intransitive priority orderings. He instead adopts the approach of utilising weak priority orderings that provide a relational comparison, as opposed to giving an objective or quantifiable value (even a relational value such as percentages). This means that the Priority Structure account can only allow for broad comparative (relational) values. For example, ‘1 is prioritised over 2 and 2 is prioritised over 3’. This view says nothing about the independent values of these priorities or the distance between their values. It only provides a vague ordinal structure.

As I have already discussed in section 7.5.2, intransitive priorities cannot be objectively ranked, and in order to account for them, Watzl claims that priorities themselves should not be considered to have objective values.

However, I am not convinced that adopting weak priority ordering in order to allow for instances of intransitive priority ordering is a good idea for two main reasons. First, I am not convinced about the existence of intransitive priorities. Second, even if intransitive priorities do exist, I do not think that they are worth accounting for at the cost that the Priority Structure account would need to pay in order to account for them.

Watzl (2017) claims that it would not be useful to assign independent values to priorities, since the resulting order of the priority structure is the same regardless of whether or not the priorities are assigned values.

The absolutist would solve this problem by first assigning a priority value to each book (maybe a number in the interval $[0,1]$), and then use that value to order the books accordingly, e.g. from the one with the lowest priority value to the one with the highest value. Since the point of prioritizing is to arrive at the final ordering, this procedure is inefficient. It uses a lot of information in the first step (is the value of the book 0.4 or 0.5?) that is then discarded at the second step (since the exact priority value will make no difference to the ordering). The same holds for attentional prioritization: the computational problem that such prioritization solves (organizing the mind) is solved most efficiently without absolute levels of priority. So, we have no reason to think that our minds evolved to operate with absolute priorities. (Watzl, 2017, p.97)

However, I would disagree with this claim that it is not useful to assign values to priorities. I would argue that not assigning values to priorities causes the following difficulties for the Priority Structure account. First, it causes difficulty with assessing the distance between priorities and, as a result, difficulty with assessing the distance between degrees of attention. For example, two different subjects could have the same priority structure (A more than B more than C) but subject 1 could have placed a significantly greater degree of priority on A than on B or C, while subject 2 could have a fairly even distribution of priority between A, B, and C (where the difference in the degree of priority between A, B, and C is minimal). Utilising weak priority ordering makes it difficult for the Priority Structure account to account for the difference between these two cases (since both have the same weak priority structure: A more than B more than C).

Second, it causes difficulties with comparing instances of attention or prioritisation that are not simultaneously co-occurring with one another. The Priority Structure account struggles to account for comparison between past and future instances of attention or prioritisation. This is because these instances are not comparable using weak priority relations, which compare active priorities in a relational way. Only occurrent states can be prioritised on this account. For example, if I have the priority structure 'A more than B more than C' at time T1 and 'D more than E more than F' at time T2 the Priority Structure account (with weak priority ordering) would be unable to account for a difference in the level of prioritisation between 'B' and 'F' since these two priorities were never relationally compared (since they were never together in a priority structure). However, it seems like we should want to be able to account for degrees of priority between different things, even if they were never put in direct comparison at the time of their prioritisation. For example, I might want to say that I prioritised my efforts to avoid being hit by a car at time T1 more than my efforts to pour milk into my tea at time T2, but since these two things never occurred together within the same priority structure, they are not comparable on this version of the Priority Structure account. They lack their own independent values that could be used to make comparisons against other priorities that were not within the same priority structure at the same time.

Section 7.8 – Wu’s Criticisms of the Priority Structure Account

In this section, I will defend the Priority Structure account from criticisms that have been made by Wu. Wu (2019) makes two main criticisms of the Priority Structure account in his review of Watzl’s book ‘Structuring Mind’. First, Wu criticises the Priority Structure account for not being supported by current empirical methodology in the same way that he claims his Selection for Action (chapter 5) account is. Second, Wu claims that the Priority Structure account does not account for cases of appropriate response in instances of psychological saliency.

Wu (2019) encourages us to:

Note that selection for action, not priority structuring, already has a pervasive foothold in experimental work on attention. In conducting the type of experiments that Watzl frequently cites, scientists maintain control of how their subjects attend by designing experimental tasks such that correct performance necessitates that subjects select some X to guide performance (for example, verbally repeating one of two verbal streams, committing specific targets to memory, reporting the presence of targets in a stream of distractors, and so on). Where subjects measurably succeed in the instructed task, the experimenter is assured that the subject, in selecting X, is thereby attending to X. This methodological assumption expresses a form of selection for action. Because it is woven into experimental methodology, the selection for action account should be the default empirical account of attention. It is an advantage of that view, unlike the priority structuring view, that it is embedded in empirical work. (Wu, 2019, p.948)

It is clear that Wu is stating that his account should be held over the Priority Structure account because current empirical tests are designed in a way that Wu claims is most compatible with his own account.

However, I will argue that Wu’s objection does not hold for the following reasons. First, it is unclear why a theory being in accordance with current empirical methodology makes that theory superior. Current empirical methodology reflects our current best efforts at measuring attention but does not reflect the direct measurement of attention itself. Attention may not be what is being measured by current methodology; there are limitations to our current empirical methodology. Second, by making this claim, Wu is using ‘is-ought’ reasoning. Just because this is the current way that we are attempting to understand attention does not mean that this

is the way that attention ought to be understood or characterised. Third, it is unsurprising that Wu's Selection for Action account aligns with current empirical methodology, as Wu's account relies on actions which are exactly what are measured by current empirical studies of attention. The fact that actions provide the only way for us to currently attempt to measure attention does not mean that these actions actually reflect attention or attentional processes. Furthermore, there may be additional non-action-based forms of attention that current empirical methodology may fail to account for because they are not, currently, directly measurable.

Furthermore, I will argue that Wu is mistaken to claim that the Priority Structure account is incompatible with current empirical methodology. Forms of attention that have measurable actions or that are based on the performance of tasks can still be measured on the Priority Structure account. It is simply the case that, on this account, these are not the only forms of attention that exist. Attention may still have measurable behavioural correlates on the Priority Structure account and may still be measured using current empirical methodology.

Finally, Wu makes it clear that his account has been developed in order to explain current empirical methodology and as such has been developed with current empirical methodology in mind. It is therefore unsurprising that his account covers current empirical methodology and therefore current empirical methodology should not be used as a means to test his account. Wu has created an account that is based on current empirical methodology and, that is using current empirical methodology as a means of support. I have already argued in chapter 5 that this is cyclical reasoning.

Wu's second objection against the Priority Structure account, is that it does not account for cases of appropriate response in instances of psychological saliency. Wu (2019) provides the following example to illustrate his criticism:

Those of us unlucky enough to stay at the conference hotel at an Eastern APA meeting some years ago were awoken in the early morning to a fire alarm. (Yes, there was an actual fire.) My first act was not to reorganize my priority structures but to get my terrified child out of the building. If my auditory experience commanded anything, it was this: get out! Of course, all theories are committed to a reorganization of the mind when attention is captured, but why think that in the design of attentional capture, the point would be to command the agent to do something internal before initiating a prudent, even life-saving act? The goal of a capacity to have attention captured is to

allow the creature to respond directly to external changes and not to fiddle with its internal world. (Wu, 2019, p.947-948)

Here Wu is criticising the Priority Structure account by claiming that it does not allow for appropriate responses during instances of saliency. Wu claims that during instances of saliency, such as the fire alarm incident, adopting the Priority Structure view would entail spending time organising our internal priorities as opposed to immediately taking appropriate action.

I will argue that this is an uncharitable interpretation of Watzl's view for the following reasons. First, because instances of extreme saliency such as the fire alarm incident play an important role in the forming and changing of our priorities—saliency is something that can (and often will) guide prioritisation on this account. It seems that a change of priorities certainly occurred when Wu made the decision to quickly evacuate the building. So, it seems that Wu has reoriented his priorities in a way that informed action, whether he was consciously aware of this process of prioritisation or not. Second, this is an uncharitable interpretation because Wu seems to imply that these judgements relating to prioritisation are time-consuming. It seems plausible to assume that a condition of extreme saliency could inform a snap judgement, but that this judgement still included a swift changing of priorities. The Priority Structure account does not state that prioritising or reprioritising is something that necessarily needs to be time-consuming in the way that Wu characterises it as being. Finally, this interpretation is uncharitable because Wu wrongly interprets the Priority Structure account by suggesting that the restructuring of priorities is necessarily a conscious process that is performed by the individual. Wu is implying that the individual sits and consciously restructures their priorities, as opposed to taking action. Watzl (2017) is very clear that attentional prioritisation can be guided by a range of different guiding principles – some are not necessarily conscious (such as salience, value, and selection history) and some can be conscious (such as goals, plans, intentions, and executive control). Prioritising can be either consciously informed or unconsciously informed on this account. The individual does not need to consciously restructure their priorities in place of taking action.

Section 7.9 – Filling in Wu’s Framework

In this section, I will summarise the core claims of the Priority Structure account by briefly considering how it might be used to fill in Wu’s theoretical framework (which I have previously discussed in chapter 5).

The **metaphysical** question: what IS attention?

Attention is a meta-process (more specifically an activity) that structures the mind by prioritising mental states, mental processes, and events.

The **properties** question: what are characteristic features of attention?

The properties of attention are outlined in section 7.3.1. Attention has the following four properties that help to distinguish it from other mental processes (such as memory).

First, it concerns mental states, processes, and events that are ‘activated’ (this is to say that it only organises occurrent mental states, processes, and events).

Second, it is guided by the following guidance principles: general world knowledge and meaning, intentions, goals and plans, selection history, prior reward and associated value, emotional factors, and executive control. Additionally, it is guided by salience (which is something that may be influenced by the other guiding principles).

Third, it is meta-representational. Attention itself is not something that can be accurate or inaccurate, nor can it be satisfied or unsatisfied. However, it is important to note that attention organises other mental states that may be representational.

Finally, attention is usable. Attention allows for action selection, belief formation, and accessible conscious experience.

The **function** question: what role does attention play?

Throughout this chapter, I have shown that attention plays several roles on the Priority Structure account. The main role of attention is to organise the mind by structuring mental states, events, and processes. However, in doing so, attention is also able to perform several

other, related, roles such as allowing for flexible action and for accessible conscious experience.

The **mechanism** question: How is attention implemented?

Watzl has made it clear in his arguments against reductionism that attention cannot be reduced to any specific or broad neuronal mechanism. However, this does not mean that there cannot be any such mechanisms that are associated with certain types of attention. The Priority Structure account can allow for the identification of highly specific attentional mechanisms (for example, a mechanism associated with visual attention to square targets during a visual search task), but there is no unifying mechanism of attention on this account. No single mechanism that is associated with certain instances of attention can overcome the heterogeneity problem, and as such attention should not be identified by its associated mechanisms.

The **consciousness** question: what is the relation between attention and consciousness?

In section 7.6 of this chapter, I have outlined the relationship between attention and consciousness on the Priority Structure account. On this account, attention structures consciousness and allows for accessible conscious experience. Additionally, attention itself has a phenomenology on this account, which is determined by the structural relation between the contents of the priority structure.

It is a further strength of the Priority Structure account that it is able to provide an answer to all five of the questions on Wu's framework, as this shows that it is explanatorily complete.

Section 7.10 – Chapter Conclusion

In this chapter, I have introduced and assessed the Priority Structure account of attention. I have shown that it is well suited to addressing the heterogeneity problem of attention, as it is able to provide a broad account of attention. This is one reason why the Priority Structure account would provide a strong foundation for a revised account of Monotropism. In the next chapter, I will further demonstrate the suitability of the Priority Structure account by

considering the compatibility between the core aspects of this account and the core aspects of Monotropism. I will do this by demonstrating the explanatory power of such a revised account, which I will title 'Priority Monotropism'.

Chapter 8 – Introducing Priority Monotropism

Section 8.1 – Chapter Introduction

In chapter 2, I explained that overcoming the heterogeneity problem is the most serious challenge that any essentialist theory of autism faces. I outlined four essentialist theories of autism: Weak Central Coherence theory, Enhanced Perceptual Functioning theory, Theory of Mind theory, and Executive Dysfunction theory, and demonstrated that none of them are able to overcome the heterogeneity problem.

In chapter 3, I introduced a further essentialist theory of autism, Monotropism, which has more potential to overcome the heterogeneity problem. However, I have argued that it is insufficient in its current form, primarily because important terms, such as ‘attention’ and the related term ‘interest’, were not clearly defined on this account. I have argued that a revised account of Monotropism, which clearly specifies these terms, is required in order to overcome these problems.

In chapter 7, I introduced Watzl’s Priority Structure account of attention, which I will be using as a foundation for a revised account of Monotropism. This will allow for the creation of a new theory of autism, Priority Monotropism, which is able to overcome the problems of the extant account of Monotropism.

In this chapter, I will outline the Priority Monotropism account and demonstrate its explanatory power by showing that, unlike the other essentialist theories, it is able to effectively tackle the heterogeneity problem. In addition, I will show that Priority Monotropism can explain all of the aspects of autism that these other essentialist theories attempt to explain, thereby demonstrating that Priority Monotropism has as much explanatory power as all of these other theories combined, plus its own additional explanatory power.

Section 8.2 – Introducing Priority Monotropism

In this chapter, I will outline a revised theory of Monotropism, which I will call ‘Priority Monotropism’. This theory will be developed by using Watzl’s Priority Structure account of attention as a foundation (which was previously discussed in chapter 7). In this section, I will outline the differences between extant Monotropism and Priority Monotropism that arise

from adopting the Priority Structure account as a foundation. I will then briefly recap by outlining the core claims of Priority Monotropism.

Section 8.2.1 – Revising Monotropism: Priority Monotropism

Adopting the Priority Structure account of attention as a foundation allows Priority Monotropism to account for the core aspects of autism that are outlined by extant Monotropism (which were previously discussed in chapter 3).

However, there are some important theoretical differences between Priority Monotropism and extant Monotropism that result from adopting the Priority Structure account. These differences are advantageous to Priority Monotropism as they allow for the development of a more explanatorily complete account of Monotropism. I will discuss these differences in the following subsections.

Section 8.2.1.1 – Removing ‘Attentional Resources’ and Capacity Limitation

As previously discussed in chapter 3, extant Monotropism subscribes to the notion of capacity limitation and is reliant on the existence of general attentional resources of some kind that are limited in some way.

Capacity limitation served a functional role on the extant account, as it provided a reason for the narrowness of attention in autism and an explanation of how narrow attention related to intensity: the fewer things these resources were distributed between, the greater the share of the resources each thing received. Capacity limitation was also used to explain why multisensory processing can be difficult for autistic individuals, the explanation being that there are not enough resources to distribute across multiple sensory modalities simultaneously. However, this claim seems to clash with the way that the extant Monotropism authors wanted to frame autism. They did not want to imply that autistic individuals are necessarily ‘disabled’ or that they have fewer resources than non-autistic individuals. However, if that is the case, then it is unclear why non-autistic people have enough resources to distribute across multiple sensory modalities, but autistic people do not. Additionally, as I have shown in chapter 3, a clear explanation of what these attentional resources are and how they function is missing on the extant Monotropism account.

However, Priority Monotropism allows for an explanation of narrow attentional distribution, perceptual intensity, and differences in multisensory processing without needing to subscribe to the notion of attentional resources.

This leads to the first major difference between extant Monotropism and Priority Monotropism: extant Monotropism relies on the notion of capacity limitation while Priority Monotropism does not. This is because, as previously discussed in chapter 7, the Priority Structure account of attention does not rely on capacity limitation. Instead, the Priority Structure account makes sense of attentional limits by appealing to the notion that prioritisation is something that is only useful if it is limited. If everything is prioritised, then nothing is prioritised. Giving priority to too many things simultaneously undermines the process of prioritisation. On the Priority Structure account, we do not need to appeal to the notion of attentional resources in order to understand attentional distribution and its limitations; these aspects of attention are explained by the fact that attention necessarily involves prioritisation.

Giving too many distinct items equally high priority without treating them as a single unit is likely to make those priorities useless and is likely to lead to interference. Consider trying to give equal priority to writing a book, caring for your children, making your love life work, training for a marathon, and learning how to play the guitar. As you add more and more items to the top of your list, you may as well stop prioritizing at all. Prioritizing is often useful only if it is limited and provides a clear ordering. It must *organize, integrate, and coordinate*. [...] If attention is prioritizing, we would expect that it will often look limited. (Watzl, 2017, p.107-108)

Priority Monotropism would provide an alternative account to that which was outlined by extant Monotropism. On this alternative account, the limitation of attention in autism is due to a narrower range of active priorities (a narrower, spiky, priority structure) during monotropic processing, as opposed to fewer attentional resources.

Section 8.2.1.2 – Differences in the Function of Attention

The second major difference between extant Monotropism and Priority Monotropism is that they both make different claims about the function of attention.

On the extant Monotropism account, it was suggested that the function of attention was to protect from overflow. This is an idea that resulted from extant Monotropism relying on capacity limitation. The idea was that these attentional resources protect the individual from being bombarded with information. On this account, attention is the process that distributes these resources in order to prevent bombardment. In contrast, Priority Monotropism takes the function of attention from Watzl's Priority Structure account, which claims that the function of attention is to organise the mind and to allow for flexible action.

The second question is: why do subjects have priority structures? An information pruning conception (according to which attention protects an organism from information overflow), and an organizational conception (according to which attention organizes the mind to allow for flexible action) are contrasted. [The Priority Structure account] argues for the organizational conception. Prioritizing plays an important role for an organism's capacity to decouple behavior from the environment. (Watzl, 2017, p.92)

As I will demonstrate throughout this chapter, this difference in the function of attention is something that allows Priority Monotropism to account for a wider range of symptoms and behaviours in autism.

Section 8.2.1.3 – Redefining 'Interests'

As previously discussed in chapter 3, the authors of the extant Monotropism account did not provide a clear definition for the term 'interests', despite this being an important term for the account. The authors of the extant account claimed that they would be using the term in a way that 'broadly coincides with common parlance'. However, as I have shown in chapter 3, this term has several common parlance uses. I have explained that they have used the term 'interest' interchangeably and without distinction, in the following three ways: 'perceptual interests/stimuli' – to describe the target of attention, 'personal interests' – to describe hobbies or special interests in autism – and 'appeal' – to describe the attraction of attention.

Priority Monotropism can bring these three types of interest together by explaining how they are linked. On this account, 'perceptual interests' refers to the perceptual processes and stimuli that are actively organised within the priority structure. However, it is important to note that on this account this type of interest is not restricted to perception. As I have shown

in chapter 7, on the Priority Structure account of attention many different cognitive processes, events, and states can be contained within the priority structure. As such, this type of interest is better thought of as ‘active interest’ or, to be more in line with the Priority Structure account, ‘active priorities’. On the Priority Monotropism account, ‘personal interests’ (or ‘personal priorities’) are related to the guidance processes that guide the organisation of the ‘active priorities’ within the priority structure. These personal interests can continue to be held even when active priorities relating to them change. For example, my personal interest in ice hockey is retained even when my active priorities relating to it change, such as turning off the TV midgame. This relates to the third type of interest used by the extant account – ‘appeal’. Things that have strong values in the range of guidance states (which I have previously discussed in chapter 7) have more ‘appeal’; they are more likely to be ‘interesting’ and guide attention. For example, special interests have a lot of associated value, and since associated value is one of the guidance principles, special interests have more ‘appeal’, and as such are more likely to be attended to.

I will further explain the interaction between these types of interest and how they can lead to the cyclical guidance of attention in autism in section 8.3.1.1.

Section 8.2.1.4 – Redefining ‘Narrow’ and ‘Broad’

The third major difference between extant Monotropism and Priority Monotropism is that Priority Monotropism provides a clearer definition for the terms ‘narrow’ and ‘wide’.

As previously discussed in chapter 3, the authors of the extant Monotropism account did not provide a clear definition of the terms ‘wide’ and ‘narrow’. However, they did provide a spatial analogy for understanding the differences between monotropic (autistic) and polytropic (non-autistic) attentional styles. On this analogy, attention during monotropic processing is comparable to a torch beam (narrow and intense), and attention during polytropic processing is comparable to a lantern (wide and diffuse). The extant authors also spoke of ‘attention tunnels’ and focused largely on the perceptual effects of attention in autism, showing how these perceptual differences led to social and cognitive differences that are also associated with autism. While they were able to successfully demonstrate the potential explanatory power of considering perceptual differences as a core aspect of autism, there are still some aspects of autism that are difficult to explain by perceptual differences alone, such as restricted repetitive behaviours and interests (RRBIs).

Priority Monotropism is able to reframe the terms ‘wide’ and ‘narrow’ in a non-spatial way. On Priority Monotropism, ‘wide’ and ‘narrow’ refer to the shape of the priority structure and therefore to the range of active priorities that will be prioritised within the priority structure. On this account, there is a narrower range of active priorities (fewer priorities) during monotropic processing and a wider range of priorities (more priorities) during polytropic processing. It is important to note that these priorities may be related to perception but are not necessarily so. The Priority Structure account allows for the prioritisation of a wide range of mental states and processes. So, while I would agree that many of the non-perceptual symptoms and behaviours in autism are best explained by perceptual differences on the Priority Monotropism account, it is not necessary for all of the non-perceptual aspects to be explained by perceptual aspects. Priority Monotropism allows for monotropic processing to affect more than perception directly; cognition may also be affected directly (not just affected indirectly through differences in perception).

Additionally, Watzl (2017) accounts for ‘wide’ and ‘narrow’ distribution through specific attentional distribution styles, in much the same way that extant Monotropism accounts for ‘wide’ attentional distribution as being ‘polytropic’ and ‘narrow’ attentional distribution as being ‘monotropic’. Watzl (2017) describes these distribution styles as follows:

Cases where attention is narrowly directed at a specific object, property, or location in space are treated as a special case, a specific distribution of attention. Narrow focus corresponds to what I will call *spiky* priority structures, where there is a single priority and everything else gets equally deprioritized. Diffuse attention, by contrast, is characterized by less spiky priority structures. (Watzl, 2017, p.73)

On the Priority Monotropism account, monotropic processing is a narrow/spiky attentional distribution style while polytropic processing is a wide/diffuse attentional distribution style. The addition of spiky and diffuse attentional styles, which is introduced by taking the Priority Structure account as a foundation, allows for an improved explanation of why attention is often narrowly distributed during monotropic processing. This attentional style is characterised by the prioritisation of only a few priorities and the deprioritisation of everything else, which, on the Priority Structure account, just is to attend narrowly.

Section 8.2.1.5 – Differences in Understanding Inattentional Blindness

The fourth major difference between extant Monotropism and Priority Monotropism concerns the relationship between attention and consciousness.

While no explicit claims are made about this relationship on the extant Monotropism account, it seems to be implied that attention is both necessary and sufficient for consciousness.

Inattentional blindness is a core aspect of extant Monotropism, and the extant authors frame empirical evidence and self-reported experiences concerning inattentional blindness in a way that implies that attention causes conscious experience and that inattention causes a lack of conscious experience. They characterise autism as a condition in which a lack of widespread attention frequently leads to widespread inattentional blindness.

Priority Monotropism does not need to be committal about the relationship between attention and consciousness. While inattentional blindness is still an important aspect of the Priority Monotropism account, it is non-committal about whether the inattentional blindness needs to be at the level of phenomenal consciousness, access consciousness, or both. That is to say that Priority Monotropism is not committed to explaining whether there is a genuine lack of phenomenology during instances of inattention or whether there is simply a lack of awareness about such phenomenology. On this account, it may be the case that autistic individuals do have phenomenal consciousness of stimuli that they are not attending to, but that these experiences are not accessible. All that matters is that there is some form of reduced awareness during instances of inattention that can explain common symptoms and behaviours in autism (such as not responding to one's name being called).

Section 8.2.1.6 – The Phenomenology of Attention

While extant Monotropism discusses how monotropic processing leads to attentional differences that can cause differences in perceptual phenomenology in autism, it does not discuss the phenomenology of attention itself.

As previously discussed in chapter 7, attention itself has a phenomenology on the Priority Structure account. This consists of the phenomenology of not just the contents of priorities within the priority structure but also of the phenomenology of the relational structure between these priorities. Each priority structure has its own phenomenology based on both the

contents and the relational structure. This phenomenology forms the perspective of the subject. There is something it is like to attend separately to A, B, and C, and there is something it is like to attend to A more than B more than C, which is phenomenologically different from attending to B more than A more than C (even if the phenomenology of A, B, and C themselves remain fixed). This is because there is something that it is like to have that specific relational structure separately and in addition to what it is like to have an experience of the contents within that structure.

This means that on the Priority Monotropism account there is something it is like to have an autistic perspective, as there is something it is like to utilise monotropic processing (regardless of the specific content that is attended to). This explains why many autistic people claim that their autism is not something that is separable from themselves as an additional condition, but rather that it is an inseparable part of their identity. On Priority Monotropism, autism is something that influences an individual's perspective through the formation of priority structures with distinct, spiky, relational structures. This further expands the account given by extant Monotropism by stating that there is a phenomenology of monotropic processing itself, which is common to autism. Despite this being an important aspect of autism and the autistic experience, it is an aspect of the condition that has been difficult for many autism theories (including the extant account of Monotropism) to account for.

Section 8.2.2 – Summarising the Core Claims of Priority Monotropism

The core claims of Priority Monotropism are as follows:

Priority Monotropism – Monotropic processing is the default processing style for autistic individuals. Utilising monotropic processing means that the individual will form spiky, as opposed to diffuse, priority structures. Within these spiky priority structures, fewer priorities are simultaneously prioritised while everything else is deprioritised. This is what it is to have a 'narrow' distribution of attention on this account.

Priority Monotropism can account for a wide range of symptoms, behaviours, and differences that are common to autism, since many different mental states can be priorities, not just perceptual states. In the next section, I will examine how Priority Monotropism can tackle the heterogeneity problem of autism.

Section 8.3 – Resolving the Heterogeneity Problem

As discussed in the previous section, the fact that Priority Monotropism moves away from providing an account of autism that focuses on perceptual attention to providing an account based on attention more generally allows for Priority Monotropism to explain a wider range of symptoms and behaviours. This move has been possible due to the broad account of attention that the Priority Structure account has provided, which allows for terms like ‘narrow’ and ‘wide’ to explain a much wider range of mental states (not just perceptual states). On the Priority Monotropism account, narrowness refers to a spiky structural relation within priority structures, and since attention organises many different mental states and processes on the Priority Structure account, Priority Monotropism can account for a wide range of mental states and processes.

“*Perceptual* attention [...] organizes the subject’s perceptual states; attention *to one’s bodily sensations* organizes the field of sensations; and *intellectual* attention organizes one’s trains of thought.” (Watzl, 2017, p.71).

On the Priority Monotropism account, autism should not be reduced to a specific range of symptoms or behaviours. First, because it is not possible to do so, while still considering autism to be a unified condition. This is because, as I have previously discussed in chapter 2, not all autistic individuals share the same set of symptoms and behaviours. Second, it is not necessary to do so on this account, as Priority Monotropism is able to identify a core essence of autism. On this account, autism is a condition that is characterised by a default tendency to utilise a specific style of attentional distribution, monotropic processing, as opposed to a specific set of behaviours or symptoms.

However, this does not mean that Priority Monotropism disregards these behaviours and symptoms that are commonly associated with autism. It treats these behaviours and symptoms as common features of autism but allows that an autistic individual need not display all of these behaviours (or even a certain number of them) in order to be considered as autistic. These behaviours and symptoms do not themselves provide an answer to the metaphysical question ‘what is autism?’ but they do provide insight into the underlying mechanisms and the core aspect of autism. Additionally, these symptoms and behaviours may still be potentially useful for indicating autism in lieu of objective tests.

Instead, Priority Monotropism focuses on providing an essentialist account that can explain why many of these common symptoms and behaviours arise in autism.

Additionally, it is possible that future research efforts may identify specific neuronal markers, genetic markers, or causes of autism. This is not a problem for Priority Monotropism, as this account is non-committal about these markers and causes—since it provides an answer to the metaphysical question ‘what is autism?’ that does not rely on the identification of such markers or causes.

As I have previously discussed in chapter 2, it is extremely difficult to identify an essence of autism, as autism is a highly heterogeneous condition. The heterogeneity problem poses two main difficulties for an essentialist theory of autism:

First, there is a wide range of seemingly disparate symptoms and behaviours that are associated with autism. It is difficult to provide a theory of autism that can connect these seemingly unrelated symptoms and behaviours.

Second, there is individual variance in the symptom profiles of autistic individuals. Autistic individuals do not always share identical profiles of symptoms and behaviours. Furthermore, an autistic individual may not even share any symptoms or behaviours with another autistic individual.

A successful essentialist theory of autism should be able to address both of these issues by showing how the wide range of symptoms found in autism are connected and also by accounting for variance in the symptom profiles of autistic individuals.

In this section, I will show how Priority Monotropism does both. I will begin by explaining how Priority Monotropism can account for a wide range of symptoms and behaviours before going on to account for individual variance in section 8.3.2.

Section 8.3.1 – Addressing the First Prong of the Heterogeneity Problem: Accounting for a Wide Range of Seemingly Disparate Symptoms

In this section, I will explain how Priority Monotropism accounts for a wide range of symptoms and behaviours that have been difficult for other essentialist theories of autism to account for (as was previously shown in chapter 2). I will explain how Priority Monotropism accounts for the wide range of symptoms and behaviours that were outlined in chapter 2. This will allow for an improved understanding of how these symptoms and behaviours are related, showing that they are not as disparate as they may initially seem. These symptoms were grouped into six main categories: differences in behaviours and interests, communication

differences, sensory differences, differences in motor skills, stress, and difficulty with fitting in. In this section, I will only discuss the first four categories, as symptoms in the last two of these categories can be explained by considering the effects of symptoms and behaviours within the first four categories. Additionally, stress and anxiety will be further discussed in chapter 9.

Section 8.3.1.1 – Accounting for Differences in Behaviours and Interests: Restricted and Repetitive Behaviours and Interests

As I have previously discussed in chapter 2, restricted and repetitive behaviours and interests (RRBIs) are a wide category of symptoms that are common in autism. This includes hyperfixation, holding special interests, stimming behaviours, and aversion to change/insistence on sameness. RRBIs are difficult for extant Monotropism to account for, as it is difficult to explain how differences in perceptual attention alone can lead to RRBIs.

However, Priority Monotropism is able to provide a better explanation of the role that RRBIs play in autism and how RRBIs relate to attentional differences. Motivation and usefulness are important factors that play a key role in determining the order of prioritisation on the Priority Structure account of attention. Motivation and usefulness are subjective and relate to an individual's interests. An individual with limited and intense interests may be more highly motivated to prioritise information that is useful in relation to these interests over and above anything else. This can lead to a narrower distribution of attention as these interests are more highly prioritised.

Adopting the Priority Structure account as a foundation also allows Priority Monotropism to provide a further potential explanation of the role that restricted interests play in guiding attentional selection.

Watzl (2023) claims that the following states guide attention:

- 1) general world knowledge and meaning
- 2) intentions, goals and plans
- 3) selection history
- 4) prior reward and associated value
- 5) emotional factors

Restricted and repetitive interests could affect the influence of these guiding states. For example, general world knowledge and meaning could be reduced by increased attendance to restricted interests, as an individual who is motivated primarily to learn about their special interest is less likely to develop general knowledge outwith that interest. Additionally, intentions, goals and plans are also likely to be related to the area of special interest. The same can also be said for rewards and associated value, which are also likely to be related to special interests (with special interests being considered as valuable and rewarding).

Additionally, having a narrower range of initial interests could lead to having a narrower selection history to draw from in the formation of future prioritisations – leading to cyclical guidance that promotes repetitive interests. The basic idea is that during monotropic processing (which is the default style in autism on this account), autistic individuals narrowly attend by narrowly prioritising. Holding this initial narrow range of priorities affects the degree of influence of the different guiding states that typically guide attention. The range of selection history and general world knowledge may be diminished by monotropic processing, and the influence of prior rewards and associated value may be increased. This leads to the cyclical reinforcement of narrowly distributed attention, as selection history and general world knowledge are limited while things that relate to the initial set of priorities are amplified. This could potentially explain how the ‘restricted’ element of these behaviours relates to the ‘repetitive’ element and could also explain the preference for sameness that is common to autism.

Furthermore, RRBI are to be expected on the Priority Monotropism account as, according to the Priority Structure account of attention, the function of attention is to organise the mind in order to allow for flexible action (Watzl, 2017). It is plausible that differences in attention can lead to differences in flexibility for action. Priority Monotropism claims that autistic individuals are more likely to utilise monotropic processing (and as a result that they are more likely to form narrower, spiky priority structures). This narrower range of active priorities could limit the degree of flexibility during monotropic processing – leading to less flexible behaviours and to RRBI. While in typical cases attention can allow for improved flexibility by narrowing the range of available choices so as to allow for selection for action, during monotropic processing this range is further narrowed, leading to a very restricted range of choices, thus reducing flexibility. This is then cyclically reinforced by restricted guiding states, which are based on past experiences – further motivating restricted interests and restricted behaviours, which in turn, further reinforce restricted guidance and so on.

This not only explains the prevalence of RRBIs in autism, but it also allows Priority Monotropism to explain other aspects of autism, such as difficulties with adaptation and task switching. The cycle of restricted guiding states reinforcing narrow prioritisation can make it difficult for autistic individuals to change their priorities in order to effectively switch tasks, as during this cycle guidance promotes repetition and increases the difficulty of task switching.

As previously discussed in chapter 2, autistic individuals commonly hold intense levels of interest in one or a few topics. These interests are commonly known as ‘special interests’. These special interests are often highly prioritised, leading to hyperfixation where they are focused on intensely and/or for a prolonged duration. During hyperfixation everything else is deprioritised.

These experiences of special interests are discussed in the following self-reported experiences from autistic participants in a study by Long (2025):

I'd explain it as something that the autistic person's really into, hyperfixates on perhaps, and spends a lot of their spare time with, and maybe even makes a career out of. (Long, 2025, p.104)

I think it's sort of an interest of sort of high enough intensity that it sort of takes up sort of a lot of your thoughts throughout the day and week, and sort of you want to always sort of put more energy into and more engagement into sort of over a prolonged period of time. (Long, 2025, p.104)

A special interest I think is associated with a large amount of hours spent thinking about it, and is often the default headspace I fall back to when I'm not doing anything else at the moment. (Long, 2025, p.105)

I have this experience when something becomes a special interest where like, I become so focused on it. I can, like absorb so much like, like, almost all the information about it within like a few days. (Long, 2025, p.105)

The way I understand special interests is they're kind of the framework by which I understand the world around me. (Long, 2025, p.105)

Additionally, special interests are often highly valued and rewarding for autistic individuals to pursue. Long (2025) reported that autistic participants described their special interests as being important to their well-being:

Special interests are something that I'm motivated to do, even when I'm not motivated to do things. Like when I'm depressed, they give me a reason to get out of bed. ... They give me structure and often in times where I feel very hopeless or overwhelmed, or stressed or burnt out (Long, 2025, p.106)

...when you're able to grasp on to a special interest, I think it's a lot easier, mental-health-wise, and ability-to-live wise. (Long, 2025, p.106)

On the Priority Structure account, associated value or reward is something that affects psychological salience. This means that aspects relating to an individual's special interests may be more salient to that individual and, as such, will be more likely to capture their attention.

Psychological salience is also penetrated by a subject's expected reward. If a lot depends on finding red berries in the forest (your life is lost without them), they are more likely to capture your attention. Financial rewards for detecting some type of stimulus have the same effect. Importantly, what is associated with a high reward may not be the same as what is currently relevant. The berries that are strongly associated with those sweet calories may now be of no use for finding a safe place to sleep or for getting some exercise. Task-irrelevant, but previously rewarded stimuli, are often psychologically salient. Alcohol- and nicotine-related cues, for example, are likely to capture the attention of those addicted to the relevant drug. (Watzl, 2017, p. 136)

Furthermore, special interests often involve expertise, as autistic individuals often spend long periods of time engaging in and researching the topics of their special interests. Watzl (2017) claims that expertise is another aspect that affects psychological salience.

“For a fluent reader of the Cyrillic alphabet, a certain letter pops out and catches her attention while seeing the same sign has no psychological salience for someone unfamiliar with that alphabet.” (Watzl, 2017, p.136).

So, given that special interests are both highly rewarding and that autistic individuals are often experts in their areas of special interests, we should expect special interests to be more salient on the Priority Monotropism account. This idea that special interests are highly salient is compatible both with the self-reported experiences of autistic individuals and with evidence from empirical studies. A study by Grelotti et al. (2005) found that Digimon characters were more salient to an autistic participant with a special interest in Digimon than familiar faces. These results were not found for a non-autistic participant with an interest in Pokémon.

The fact that special interests are more salient means that they are more likely to be attended to. This further influences future prioritisation towards these special interests, since selection history is a guiding principle for the formation of new priorities.

Special interests play an important role in the formation and reinforcement of RRBIs, according to Priority Monotropism. This is because personal interests are part of the guiding states that influence what is prioritised and what is deprioritised—they are associated with high value and influence intentions, goals, and plans. For autistic individuals, personal interests may be something that they organise the rest of their lives around. The significance of something may be measured by its level of relevance to their special interests, and as such, priorities may be structured around special interests. This can lead to the development and sustainment of RRBIs that relate to these special interests.

Additionally, as well as being highly motivating and salient, autistic individuals often report that engaging in RRBIs, such as stimming or spending time with special interests, produces a calming effect. These behaviours are often performed in response to distress or intense emotions (Morris et al., 2025).

Priority Monotropism can also provide a potential explanation for why engaging in RRBIs may be calming for autistic individuals and why these behaviours (along with other sensory seeking behaviours) may be frequently performed in times of distress.

RRBIs, such as stimming, may serve as an easy target for the redirection of narrow attention during monotropic processing. When an RBI, such as hand flapping, is prioritised, other potentially unwanted aspects of the environment are deprioritised due to the narrow prioritisation that occurs during monotropic processing. On this account, stimming provides the individual with a way to voluntarily shut out unwanted stimuli by narrowly prioritising the stimming behaviour. Stimming provides a convenient and accessible way for the

individual to deprioritise unwanted inputs. It is especially convenient since many stimming behaviours rely solely on bodily movements, such as blinking, rocking, hand flapping, or making sounds (National Autistic Society, n.d. i), which are almost always accessible.

This idea of stimming functioning to deliberately redirect attention in order to block out unwanted inputs is also expressed by Morris et al. (2025), who state that:

Autistic adults have described stimming as creating a singular, controllable point of focus in contrast to a broader uncontrollable environment. Stimming, then, has been reported to aid in relieving overwhelming sensory input by both drawing attention away from it and quelling any resulting distress or anxiety. (Morris et al., 2025, p.2)

This notion is consistent with the following self-reported experiences from autistic individuals that were reported in a study by MacLennan, O'Brien, and Tavassoli (2022):

I seek comfort in small things (such as nice textures, I currently have a coin that I keep in my pocket to hold when I'm nervous) so that I can filter the sensory environment. It helps distract me from what is happening around me so that when I seek that input, I can avoid more intense or unpleasant inputs. (MacLennan, O'Brien, and Tavassoli, 2022, p.3069)

I really like music and listen to it as much as I can during the day including during work and at all times on public transport and when shopping etc. I like being able to have full control of what I hear and how loud it is. I'm very specific in the music I listen to and will sometimes just listen to the same band/album/song for weeks... (MacLennan, O'Brien, and Tavassoli, 2022, p.3070)

Furthermore, understanding the importance of special interests on the Priority Monotropism account may allow us to explain why autistic individuals often focus their attention on different targets than non-autistic individuals. Autistic individuals often have different priorities than non-autistic individuals. Special interests are prioritised over stimuli that are typically prioritised by non-autistic individuals (such as faces) (Grelotti et al., 2005). The Priority Monotropism account can account for the fact that autistic participants often attend to different targets than non-autistic participants on eye-tracking tests. When visually attending to a scene, non-autistic participants typically look at the faces of people within the scene while autistic participants typically do not (Rigby, Stoesz, and Jakobson, 2016).

This difference is explained by a difference in priorities – non-autistic individuals are more likely to prioritise the detection of faces while autistic individuals are not. These differences in priorities lead to differences in the distribution of attention.

While RRBIs include an element of ‘restriction’, it is important to note that autistic individuals are not necessarily restricted to a single special interest and may in fact have many special interests. These special interests can come in and out of rotation, where a specific special interest will hold priority over the others for a certain period of time before being overtaken by a different interest. That first interest will be deprioritised until it comes back into rotation at a later time. The case of rotating special interests leaves several questions to be answered: why do special interests rotate? How do new interests enter into the picture on the Priority Monotropism account? If the cyclical guidance loop is so strong, then how do other interests get the chance to come into rotation? Priority Monotropism provides the following answer: selection history is not the only guidance principle. When other guidance principles are sufficiently weighted, then the guidance cycle may be broken. This could be to the detriment of the individual (as is the case when salience forcibly disrupts the cycle and causes overwhelm or meltdowns) or it may not (as is the case when interests rotate normally). Interests may rotate when the other guidance principles change, for example when the ‘value’ guidance principle changes. Autistic individuals may overconsume their special interests and exhaust them by learning about them or engaging in them very intensely. When a lot of information or content is consumed in relation to an interest, the individual may start to run out of ways to engage in or research the interest. When interests are exhausted, their value decreases, and the value of new potential interests may increase, this is something that may slow the strong guidance cycle enough to allow a different interest to enter in. The individual may cycle between their range of special interests and return to a previously exhausted interest when new content, information, perspectives, or ways to engage in the interest are available to them once again.

Difficulties with adaptation, task switching and experiencing sticky attention are further aspects of RRBIs that Priority Monotropism can account for. As I have previously discussed in chapter 2, autistic individuals often struggle with behavioural flexibility and situational adaptation. This is an issue that is acknowledged by the extant Monotropism account, but which is not explained by it.

This type of sticky attention is demonstrated in the following self-reported experience of an autistic participant in a study by MacLennan, O'Brien, and Tavassoli (2022) "I can easily be completely distracted by certain decorative lighting; I can regularly spend long periods of time in the lighting section of a department store and find it difficult to leave." (MacLennan, O'Brien, and Tavassoli, 2022, p.3070)

Priority Monotropism can provide an explanation for why autistic individuals may face difficulties relating to adaptation and task switching. As previously discussed in this section, narrow attentional styles may lead to the stronger reinforcement of fewer guiding states, and these restricted guiding states may further reinforce the narrow attentional style. It may be difficult for autistic individuals to break this cycle in order to switch tasks or adapt to new situations. This can explain why attention is commonly reported as sticky in autism and why autistic individuals often have difficulty with executive control – since they must fight against a strong cycle of guidance and reinforcement in order to deliberately attend to something else. This also provides an explanation for why autistic individuals often have a preference for sameness or an aversion to change, as the narrower guiding system that is developed through the utilisation of monotropic processing promotes sameness.

Differences in executive control can be further explained by Priority Monotropism by considering the role of executive functioning and passive attention guidance on the Priority Structure account of attention. On the Priority Structure account, attention will be passively guided by default unless the subject overrides this by utilising executive control.

"Guided by psychological salience, it is argued, priority structures tend to evolve into basins of attraction. A subject's attending naturally unfolds according to psychological saliences, unless she applies executive control to structure her unfolding activity" (Watzl, 2017, p.114).

On the Priority Monotropism account, autistic individuals often have reinforced rigid guiding states that relate heavily to their special interests. As I have previously discussed in this section, these special interests are especially salient due to their associated value. Having such highly salient guiding states means that autistic individuals are likely to have priority structures that evolve into very deep basins of attraction. As a result, it may be more difficult for them to effectively utilise executive control to override these passive guiding states.

Additionally, Watzl (2017) claims that:

Active attention guidance is effortful especially when goals and action plans are novel and have not been routinized [...] Bodily activities too require more effort when they have not been routinized. And it requires special effort to work against one's well-established routines (psychologists often call this *task set inertia*). (Watzl, 2017, p.148)

Since active attention guidance (executive control) may already be difficult for autistic individuals, active guidance towards novel things, which have not been routinised, will be even more difficult. This may further strengthen the preference for sameness in autism.

Another aspect of autism that relates to RRBIs is hyperfixation. As I have previously discussed in chapter 2, autistic individuals commonly experience hyperfixation, where they narrowly focus their attention for extended periods of time. This is often referred to as 'attention tunnelling'. While the extant Monotropism account focused on narrowly distributed (perceptual) attention in autism in order to account for attention tunnelling and hyperfixation, it did not account for the temporal extension aspect of hyperfixation.

Priority Monotropism accounts for both the narrow distribution and the extended duration of attention during hyperfixation. This is once again due to the revisions that have been made as a result of adopting the Priority Structure account of attention as the foundation for Priority Monotropism. Watzl (2017) associates spiky priority structures with vigilance and the extended duration of attention:

... low arousal states tend to be characterized by distributed and not spiky priority structures. [...] levels of vigilance can be distinguished in terms of a subject's capacity to sustain attention for an extended period of time. A highly vigilant subject normally has her priorities stably dictated by the activity she is performing, while this will not be the case for the low vigilance subject. (Watzl, 2017, p.98)

This shows that spiky priority structures are related to higher levels of vigilance and increased duration of attentional extension (compared to diffuse priority structures). This allows Priority Monotropism to account for vigilance and hyperfixation, as monotropic processing is characterised by spiky priority structures on this account.

Section 8.3.1.2 – Accounting for Communication Differences

Section 8.3.1.2.1 – Differences in Language Use and Language Development

As I have previously shown in chapter 2, there is a wide range of common symptoms and behaviours in autism that relate to differences in communication and language use. These differences can make social communication more difficult for autistic individuals.

Whilst this issue is not explicitly discussed by Watzl (as he is aiming to make an account of attention in general as opposed to an account of attention in autism), differences in language development could plausibly be explained by the Priority Structure account of attention. This is because Watzl (2017) claims that, on the Priority Structure account of attention, attention is an important aspect of successful language development. Therefore, differences in attentional distribution could plausibly lead to differences in language development. This is a further point of compatibility between the Priority Structure account and Monotropism.

Similarly to extant Monotropism, Priority Monotropism explains most of these differences in communication and language development as resulting from differences in the distribution of perceptual attention. Narrow attentional distribution may lead to important cues during language learning being missed out on or the formation of idiosyncratic labelling (as was previously discussed in chapter 3).

However, this explanation alone does not account for the fact that some autistic individuals do not use language (or that they only use language very sparingly) despite the fact that they have acquired language skills and are able to use language.

Priority Monotropism is able to explain this aspect of autism. Autistic individuals may not be motivated to utilise language when language utilisation is not relevant to their interests. This is because doing so could be a distractor that may require them to disengage from their current interest. Since autistic individuals have a narrower range of active interests (priorities) on this account, they may engage in social communication less frequently. Many autistic individuals who use language very sparingly only do so when it allows them to bring about desired outcomes (for example, for practical reasons, such as requesting an item).

This may also explain language regression in autism; in infancy, the individual may be very interdependent on others, and they may not have the capacity to hold distinct personal interests or priorities. As the child grows, they gain more independence and gain access to

more of their environment. At this stage, there are more opportunities for them to develop their own interests and priorities and, as a result, they may engage less in communication with others. This leads to a regression in language use and language development. This is a better explanation of language regression than that which was proposed by the extant Monotropism account (which I have previously discussed in chapter 3). The extant explanation was that language may be used to manipulate the attention of the child by using commanding statements, such as “look at that cat!”. However, as I have already discussed, this explanation goes against core claims of the extant account. If autistic individuals are often unaware of many things during monotropic processing (including social cues), then it is not clear how such use of language could be used to distract or manipulate them.

Additionally, as I have previously shown in chapter 2, the development of atypical accents is common in autism. These atypical accents vary from those that are commonly associated with the individual’s local region.

One explanation is that these accents are developed as a result of less time being spent attending to social interactions with peers and more time being spent attending to media relating to special interests. A study by Beccaria, Gagliardi, and Kissine (2024) reported that these accents are acquired through TV and media consumption.

On the Priority Monotropism account, due to the self-reinforcing guidance cycle (which was discussed in the previous subsection), autistic individuals may be more likely to engage with a narrower range of interests and interest-related media that they are more likely to repetitively consume. This could lead to stronger reinforcement of learned accents, which have been acquired through the consumption of this media, as there are fewer accessible alternative sources for them to draw from.

Section 8.3.1.2.1 – Differences in Joint Attention and Joint Action

As I have previously discussed in chapter 2, autistic individuals are less likely to utilise joint attention and joint action. For example, autistic children are less likely to respond to and less likely to utilise joint attentional cues, such as pointing (Ramos-Cabo, Vulchanov, and Vulchanova, 2021). They may also struggle to engage in collaborative play with other children (Khatab et al., 2024).

Once again, this is an aspect of autism that, while not discussed by Watzl, can be explained by the adoption of the Priority Structure account of attention as the foundation for Priority Monotropism. Watzl claims that attention is necessary for social co-ordination and co-operation and that it allows agents to “link their minds and generate joint experiences and joint actions” (Watzl, 2023, p.6). Therefore, it is plausible that differences in the distribution of attention could lead to differences in joint action and joint attention on this account.

Priority Monotropism provides several different explanations for differences in joint attention and joint action in autism. First, on this account, autistic individuals may not be aware of the joint attentional cues of others due to their utilisation of monotropic processing. For example, they may be unaware of someone calling their name, pointing for them to look at something, or inviting them to play. Second, even if the individual is aware of the joint attentional cue, it may be difficult for them to disengage from their current interest in order to change to the new target. Additionally, they may not be motivated to do so, as these new potential priorities may be outweighed due to strong restricted cyclical guidance, which leads to greater prioritisation of special interests and past priorities.

Section 8.3.1.3 – Accounting for Perceptual Differences

Section 8.3.1.3.1 – Attentional Exclusion and Inattentional Blindness

As previously discussed in chapters 2 and 3, autistic individuals often report experiences of and exhibit signs of inattentional blindness. This is something that the extant Monotropism account has emphasised. However, despite placing importance on this aspect of autism, the extant Monotropism account was unable to provide an explanatory picture of how or why attentional exclusion occurs in autism. This is because, as I have previously argued in chapter 3, extant Monotropism provided a vague account that did not specify what was meant by ‘attention’ and did not specify the relationship between attention and consciousness. These claims are important for providing an account that allows for inattentional blindness.

Priority Monotropism is able to provide an improved explanation for why autistic individuals often report experiencing and show signs of inattentional blindness. As I have previously discussed in chapter 7, on the Priority Structure account, attention is necessary for accessible phenomenal consciousness. This allows for inattentional blindness to occur in cases where inattention leads to a lack of accessible phenomenal consciousness.

Another important addition to this explanation is that, on this account, priority ordering has an exclusionary element. The process of prioritising involves determining what will be included and what will be excluded from the priority structure.

Part of the placement structure, and an important one at that, surely concerns which stories are placed in the newspaper at all. Similarly, one important aspect of a subject's attentional priority state concerns not just the priority ordering of her current mental life, but also what is and what is not included. Think of our jazz-loving subject. While she was listening to the band, she did not hear at all what her friend sitting next to her told her (let's assume that she really did not hear it). Hearing her friend is not part of her (auditory) priority structure, because she is not hearing her friend at all. Part of maintaining her current priority structure consists in keeping out hearing those other things. (Watzl, 2017, p.78)

On the Priority Monotropism account, because autistic individuals are more likely to have narrow, spiky, priority structures, this means that there will be more exclusion, as they will have fewer active priorities at any given time. This exclusion can lead to the type of inattentional blindness that the extant account of Monotropism had intended to account for.

Additionally, Watzl (2017) provides evidence from Yantis and Jonides (1990) that attentional capture by abrupt onset stimuli is less effective when the subject is under an unrelated high attentional load.

They found that attention capture by abrupt onsets is not fully automatic in that sense. While abrupt onsets capture attention when the subject is not already engaged in a particular task and hence “in diffuse attention mode”, they do *not* capture attention “when other task demands have led to highly focused attention”. (Watzl, 2017, p.136)

This helps to explain why autistic individuals are more prone to experiences of inattentional blindness, even in the face of sudden onset or seemingly salient stimuli, such as one's own name being called. On Priority Monotropism, autistic individuals do not typically utilise a diffuse (polytropic) attentional style; they instead typically utilise a narrow, spiky (monotropic) attentional style. This makes them less susceptible to attentional capture by intruding stimuli (even those that may typically be salient).

Section 8.3.1.3.2 – Differences in Multisensory Processing

As previously discussed in chapter 2, autistic individuals often utilise unimodal processing over multimodal processing.

Priority Monotropism is better able to account for differences in multi-sensory processing than extant Monotropism. Extant Monotropism suggests that these differences in multisensory processing are due to differences in perceptual attention. Wherein attending intensely to one sensory modality causes exclusion in other sensory modalities.

However, what was missing on the extant Monotropism account is an explanation of how attending intensely causes exclusion. Additionally, the extant account of Monotropism does not provide an explanation of why exclusion can occur intramodally and/or intermodally. For example, why does monotropic processing sometimes make it the case that there is only attention to vision but not to the other senses (intermodal exclusion) and at other times makes it the case that there is only attention to some specific aspects of vision but not to all other aspects of vision (intramodal exclusion)?

Priority Monotropism provides an explanation for this that is similar to that proposed by the extant Monotropism account, but which is more explanatorily complete. On the Priority Monotropism account, autistic individuals often utilise a narrow attentional style (monotropic processing) wherein fewer mental states or processes are simultaneously actively prioritised. This attentional style means that mental states and processes, such as inputs from sensory modalities (for example, vision or audition), that are not prioritised will be excluded from the priority structure, and as such the individual will not have an accessible phenomenal conscious experience of them. This also allows for an explanation of intramodal exclusion, as finer-grained mental states and processes, such as shape or colour perception, may be included or excluded within the priority structure.

Additionally, Priority Monotropism can also allow for cases of sensory exclusion that are not just related to narrow perceptual attention, but rather to any type of narrow attention. For example, Priority Monotropism could account for sensory exclusion caused by cognitive prioritisation (wherein narrowly directing attention to a thought could cause sensory exclusion). This is because there are a wide range of mental states and processes that can be included or excluded from the priority structure. This allows Priority Monotropism to account for a wider range of attentional exclusion than the extant Monotropism account.

Section 8.3.1.3.3 – Simultaneously Co-Occurring Hyper- and Hypo-Sensory Sensitivities

As previously discussed in chapter 2, autistic individuals often report experiencing hyper- and hypo-sensory sensitivities. These hyper- and hypo-sensory sensitivities can co-occur within the same individual at the same time, leading to situations where the individual is both hyper- and hypo-sensitive to sensory inputs (as was demonstrated in chapter 3). This is an aspect of autism that may initially seem paradoxical.

However, Priority Monotropism can address this issue. It provides a similar answer to the extant Monotropism account as to why we often find co-occurring hyper- and hypo-sensory sensitivities in autism (outlined in chapter 3). According to both accounts, hyper- and hypo-sensitivities occur simultaneously due to narrow attentional distribution, which leads to hyper-sensitivities, and widespread attentional exclusion, which leads to hypo-sensitivities. However, Priority Monotropism provides a more detailed explanation as to the nature of this narrow attentional distribution and how it relates to conscious experience. As I have previously discussed in chapter 7, on the Priority Structure account of attention, the level of attentional priority that something has affects the likelihood that the subject will have accessible conscious awareness of that thing. This allows Priority Monotropism to explain hyper- and hypo- sensitivities in autism, as utilising monotropic processing leads to narrow, spiky prioritisation in which only a few things are intensely prioritised (resulting in hyper-sensitivities) while everything else is deprioritised (resulting in hypo-sensitivities).

Section 8.3.1.3.4 – Differences in the Perception of Global Motion Coherence

Another sensory aspect of autism that Priority Monotropism can account for is differences in the perception of global motion coherence. A meta-analysis by Van der Hallen et al. (2019) has shown that autistic individuals display reduced global motion coherence compared to non-autistic individuals. Furlan and Smith (2016) define global motion coherence as “...the ability to extract the central direction tendency from an extended area of visual space containing widely disparate local directions.” (Furlan and Smith, 2016, p.7314). For example, the ability to tell what general direction a whole group of dots is moving in, even when individual dots within that group are moving in different directions.

Adopting the Priority Structure account of attention as the foundation for Priority Monotropism allows Priority Monotropism to explain these differences. Watzl claims that: “... attention sometimes degrades the discriminability of apparent motion” (Watzl, 2017, p.169-170).

This shows that, on the Priority Structure account, attention may sometimes be responsible for reduced discrimination of motion. Since, on the Priority Monotropism account, autistic individuals utilise attention more narrowly and intensely, degradation of the discriminability of global motion should be expected.

On the Priority Monotropism account, utilising monotropic processing makes it difficult to simultaneously attend to all of the relevant perceptual cues that are required to gain an understanding of the direction of global movement.

Section 8.3.1.3.5 – Improved Detail Perception

As discussed in chapter 2, another perceptual aspect of autism is improved detail perception. This is typically evidenced by improved performance on local-over-global processing tasks, such as the Navon task (Falter-Wagner et al., 2022), Embedded Figures task (Cribb et al., 2016), and visual search tasks (Kaldy et al., 2017) (which were previously discussed in chapter 2). Successful performance on these tasks requires the participant to disregard global distractors and gestalts in order to identify specific details of the scene.

This is a further perceptual aspect of autism that can be explained by Priority Monotropism. According to Watzl (2017), one of the roles of attention is to help with decoupling sensory information from the environment. On the Priority Monotropism account, the narrow and intense distribution of attention during monotropic processing could potentially lead to greater levels of sensory decoupling. This would explain why autistic individuals perform better on these tasks, as increased sensory decoupling makes detail discrimination easier and gestalt discrimination more difficult.

Additionally, as I have previously discussed in chapter 7, on the Priority Structure account, the degree of attentional priority that something receives can affect the likelihood of the conscious experience of that thing being phenomenally accessible to the subject. Watzl (2017) states that this also affects the pop-out effect. This provides a further potential explanation for why autistic individuals typically perform better on these tasks.

Section 8.3.1.3.6 – Sensory Distress and Meltdowns

In chapter 3, I have shown that the authors of the extant Monotropism account have attempted to account for meltdowns and distress in response to loud or abruptly salient stimuli. They claimed that meltdowns are caused by the individual being pulled out of the attention tunnel, as this experience is deeply distressing for the individual (Murray, Lesser, and Lawson, 2005). However, as I have previously argued, in chapter 3, it is unclear what it means for the individual to be ‘pulled out of the attention tunnel’ or why this experience should be distressing.

Priority Monotropism provides an answer to these questions. Watzl (2017) claims that for salient cues:

... the role of the cue is such that the subject’s perception of that cue acts as a guiding state. The perception of the cue guides the subject’s attention so that a perception of the target (in a valid trial) ends up higher in the subject’s priority structure than it would otherwise have been. (Watzl, 2017, p.118)

This shows that salient cues can guide attention. As previously discussed in this chapter, on the Priority Monotropism account, monotropic processing can lead to narrower, cyclical, attentional guidance. If a competing cue is salient enough, then it may interrupt this guidance cycle by providing competing guidance. The individual may be torn between the two competing guiding states and may experience distress or ‘shut down’ as a result of not knowing which of the competing states to attend to.

This also allows for an explanation of why abrupt or salient sensory stimuli, such as loud sounds, are sometimes distressing for an autistic individual and why they are sometimes welcomed or even sought out by that same individual as a form of sensory seeking. This is another aspect of autism that is seemingly paradoxical and is difficult to account for, as it cannot be explained by blanket sensory sensitivities. On the Priority Monotropism account, it depends whether the individual is willingly and deliberately attending to the sound or whether the sound is an unwelcome interruption that is affecting their ability to focus elsewhere. In these cases, the distress is not necessarily the result of a sensitivity to the sound itself, but rather the result of a conflict in attentional guidance.

Additionally, the Priority Structure account of attention claims that there is a phenomenal difference between instances where attention is actively directed versus instances where it is captured or pulled.

There seems to be a clear phenomenal difference between a *passive* attention shift, where you feel your attention being captured by, or drawn to, a certain object [...] and a situation where you *actively* direct your attention to an object. (Watzl, 2017, p.233)

On the Priority Monotropism account, this phenomenal experience of attention being pulled away may add to the sense of distress that autistic individuals experience when they are interrupted.

Furthermore, Priority Monotropism may shed light on the question of why autistic individuals may frequently experience meltdowns. The fact that autistic individuals attend to a narrow range of active priorities during monotropic processing means that there is greater potential for a wider range of stimuli to be competing (as most stimuli are not related to the current interest). Whereas, during polytropic processing these salient stimuli may simply be integrated into the current range of active priorities without much issue, as there is space within the priority structure to accommodate them.

Section 8.3.1.3.7 – Differences in Interoceptive Awareness

Another seemingly paradoxical aspect of autism is that interoceptive awareness is both heightened and dampened. Autistic individuals report both hyper- and hypo- interoceptive awareness to states such as pain, hunger, stress, and emotions. An example of these interoceptive differences can be seen in the following self-reported experience:

My partner will point out that I'm shivering before I've realised I'm feeling very cold. If my body is in pain it takes me time to realise, for example if my partner sits on my foot and it is in an awkward position, he'll notice, and I'll become aware it's painful once it has been pointed out. (MacLennan, O'Brien, and Tavassoli, 2022, p.3070)

Priority Monotropism accounts for these interoceptive differences in the same way that it accounts for other cases of sensory hyper- and hypo-sensitivities – by stating that there will

be heightened awareness of interoceptive states that are attended to and decreased awareness of interoceptive states that are unattended to. Since attention is narrowly distributed during monotropic processing, there will be more cases of decreased awareness of interoceptive states than heightened awareness of interoceptive states.

Additionally, Kalingel-Levi et al. (2022) reported that autistic individuals often relied on external information in order to determine their interoceptive state. For example, autistic participants reported relying on visual signals of injury (such as the presence of blood or visible bodily damage) in order to determine that they were in pain. These signals allowed them to make judgements about their interoceptive states in the absence of interoceptive awareness. One participant in this study reported the following:

I was sitting to the side of the bicycle, and my leg got caught in the wheel. I felt something bad had happened, but I did not know what. Then I saw my sandal was completely torn up, and my leg was stuck; only then did I understand I was in pain. (Kalingel-Levi et al., 2022, p.6)

Additionally, this study also found that autistic participants could further reduce their awareness of their pain by hiding these visual cues. One participant reported this type of experience as follows:

A large rock fell on my leg and wounded me. The wound bled. I preferred not to pay attention to the wound; ignore it and keep working. I pulled my pants down to cover the wound so I did not see and think about my pain. (Kalingel-Levi et al., 2022, p.9)

These reports are also consistent with the explanation provided by the Priority Monotropism account, as monotropic processing makes it more likely that autistic individuals will engage in unimodal (or reduced multimodal) processing (since they prioritise fewer processes simultaneously due to having narrower priority structures). This explains why the autistic participants in this study reported relying on information from within a single modality to determine what the state of the other should be. For example, using visual cues to determine that they should interoceptively be feeling pain.

Autistic participants in this study also reported being able to reduce their interoceptive awareness by distracting themselves with something else, further emphasising this point. One participant stated that:

When I am in pain, ... the best place to escape to is my phone. There are many activities I can distract myself with—movies, games, and more. It's as if it passes the pain, and you do not think about it. (Kalingel-Levi et al., 2022, p.8)

Furthermore, some participants in this study reported deliberately engaging in self-injurious behaviours in order to create a controlled distraction:

Sometimes, when I am in a lot of pain, I pinch myself somewhere else. Then the other pain distracts me from the sharp pain.

During high school, when I was in pain, I would bang my head on the wall... One time I was waiting for a doctor's appointment, and my head really hurt, a strong migraine. I came close to my parents and banged my head on the wall next to them. It calmed me. (Kalingel-Levi et al., 2022, p.9)

Another issue that is related to interoceptive differences is that autistic individuals may sometimes fail to carry out tasks related to their basic needs due to lack of awareness of their interoceptive states. For example, a participant from a study by Kinnaird et al. (2019) states that at times they are “literally forgetting to eat because I just get too busy to, I don't know, recognise the signals” Kinnaird et al., 2019, p.6).

Interestingly, differences in interoceptive awareness are also related to meltdowns. Adams et al. (2025) claim that inattention to interoceptive awareness can lead to unawareness of mounting stress until it becomes ‘too late’ and a meltdown occurs. This makes sense of the fact that meltdowns may feel unexpected and may be difficult for the individual to effectively prevent.

This is shown in the following self-reported experience from this study: “I sometimes only realise I'm getting overwhelmed when it gets quite extreme ... that can lead to ... a bigger meltdown than I would have had if I'd spotted it earlier.” (Adams et al., 2025, p.1590).

This could also allow Priority Monotropism to explain why meltdowns are more likely to occur in certain settings and under certain circumstances, such as in a busy shopping centre or at school, for example. Narrowly attending to aspects of these environments and to expected social behaviours (by masking) can contribute to meltdowns occurring more frequently in these environments. This is because narrowly attending to these aspects reduces the awareness of relevant interoceptive cues that may indicate that stress is gradually building up. The individual may not notice this stress until it has become overwhelming, resulting in a meltdown.

This idea is compatible with a further self-reported experience from this same study by Adams et al. (2025), in which the participant reports experiencing ‘shutdowns’ (which are a stress response, similar to meltdowns, except with internally directed responses such as freezing or going silent as opposed to outwardly directed responses such as screaming) in response to attending to socialisation at the expense of attending to their interoceptive state at school.

My ability to read myself lessens quite a bit because I’m so focused on other people and their actions and their experiences and I think that’s why I ended up having so many um shutdowns in school ... I wasn’t reading the signs that I was getting overwhelmed. (Adams et al., 2025, p.1592)

The idea that masking can be a distractor that reduces the individual's ability to interoceptively attend is also supported by another self-reported experience from this study:

When I was masking at school like I wouldn’t really feel hungry or thirsty and stuff and like I often wouldn’t even go to ... the toilet until I literally got home got through the front door and then like kind of unmasked more. [...] I noticed since I started homeschooling ... I’m doing my own thing that then I seem to kind of realise oh I need to go to the toilet I need to drink because I’m thirsty. (Adams et al., 2025, p.1592)

Furthermore, considering interoceptive differences may also allow for an improved understanding of social communication differences in autism. Autistic individuals may use their own experiences, as opposed to social norms, as a guide in order to navigate social

situations (Nah and Poon, 2011). For example, if an autistic person is being told about a recent negative experience of the person that they are talking to (for example, that they recently lost their job), then they may respond by sharing a story about a similar experience that they have had, as opposed to responding with statements of sympathy. This response is intended to show the speaker that they understand their situation.

Since autistic individuals may utilise their own experiences as a guide for their understanding of the experiences of others, differences in interoception (which may affect their understanding or awareness of their own experiences) may affect the way in which they relate to others. This could explain why autistic individuals may sometimes give responses that may unintentionally cause offence or be seen as inappropriate. For example, if a non-autistic person is telling an autistic person that they tripped in front of others at school (and in doing so they are implying that this was an embarrassing or upsetting experience), then they may be seeking a reassuring response. However, the autistic individual may not understand that this warrants a reassuring response, as they may not understand that this is an embarrassing event. They may have had a similar experience of tripping in front of others at school, but they may not have been interoceptively attending to their emotions at that time in order to form the relevant association of this event as an embarrassing event. This kind of understanding is important for relating to others and giving appropriate responses to the experiences of others. It is important to note that, on this view, it is not necessarily the case that the autistic individual is unsupportive or uncaring, as the inappropriate response was not intentional in this example. Rather, it may be the case that they might not always have a full understanding of the emotions or sensations that are typically associated with certain events. Whether ‘appropriate’ associations are formed would depend on whether or not they were interoceptively attending in order to form these associations.

This may also contribute to the double empathy problem, as it could also explain why non-autistic individuals may find it difficult to relate to autistic individuals. Autistic individuals may also form unusually strong or atypical associations in cases where interoceptive awareness is narrowly attended to. For example, if the individual interoceptively attends to their sensation of discomfort while also attending to a song, then they may form an association between the song and their feeling of discomfort, as opposed to forming an ‘appropriate’ association between their sensation of discomfort and the source of their discomfort (which was unattended to in this situation). This is something that may be difficult for non-autistic individuals to understand.

Section 8.3.1.4 – Differences in Motor Skills

As previously discussed in chapter 2, another common set of symptoms and behaviours in autism relates to differences in motor skills. Priority Monotropism explains these differences through differences in the distribution of attention. The successful utilisation of motor skills, such as balancing, can require the integration of a wide range of sensory inputs from across modalities such as vision, proprioception, and touch. For example, when riding a bike, one needs to be aware of the road in front of them and the space around them while also being aware of their bodily position and movements (such as making sure that they are appropriately turning the handlebars and pedalling). Utilising monotropic processing can make it difficult to attend to and co-ordinate all of these elements together, whilst also attempting to attend to cognitive aspects of the task, such as planning precise movements.

Section 8.3.1.5 – Section Conclusion

In this section, I have shown that Priority Monotropism can overcome the first prong of the heterogeneity problem, as it can account for a wide range of seemingly disparate symptoms and behaviours that are common to autism. I have demonstrated that it can provide an explanatory account that shows how these symptoms and behaviours are related and explains why they often occur in autism. These symptoms and behaviours are all related to monotropic processing. As I have shown in chapter 2, this is something that other essentialist theories of autism have been unable to do.

In the next section, I will further demonstrate the explanatory power of Priority Monotropism by showing how it can address the second prong of the heterogeneity problem: individual variance.

Section 8.3.2 – Addressing the Second Prong of the Heterogeneity Problem: Individual Variance

In chapters 1 and 2, I introduced the heterogeneity problem of attention. In the last section, I have shown that the Priority Structure account of attention addresses the first prong of the heterogeneity problem by explaining how attentional differences can account for a wide range of symptoms and behaviours that are common to autism. In this section, I will address the second prong of the heterogeneity problem: individual variance. This second aspect of the heterogeneity problem also poses a serious challenge that other essentialist theories of autism

have not been able to overcome (as I have previously discussed in chapter 2). In order to address the second prong of the heterogeneity problem, an essentialist theory needs to explain why there is individual variance in autism. This is to say that it needs to show why it is the case that not all autistic individuals share the same (or even similar) symptom profiles. It also needs to provide a suitable essence that can link all of these individuals despite their lack of commonality. This is to say that it needs to provide an explanation of why we should consider all autistic individuals to be members of the same group despite their potential differences in symptom profiles. I will show that Priority Monotropism is able to do both of these things.

On the Priority Structure account of attention, Watzl makes it clear that there is heterogeneity within attention and that this is not a problem for his account but rather it is something to be expected. This is because attention is broad – it can be instantiated by a wide range of mental states and processes for different purposes.

...there is diversity in attention guidance: it can be exogenous, endogenous, controlled, automatic, value and task sensitive, and more. [...] There is also diversity in which types of representational states attention organizes: there is attention in perception, cognition, emotion, action planning, etc. [...] Finally, there is diversity in the types of usability or priority: there is use for action, prominence in conscious experience, or an ordering by what is currently relevant to the subject. (Watzl, 2023, p.6)

Similarly, on the Priority Monotropism account, heterogeneity within autism is also to be expected. This is the case for the following two reasons. First, because priorities are guided by interests that can vary based on their content and intensity. Second, because attention can be constituted by a wide range of mental states and processes. On the Priority Monotropism account, to utilise monotropic processing is to have a specific, narrow shape of priority structure. However, Priority Monotropism is not committed to any claims about the specific contents of the priority structure. These contents are likely to vary from individual to individual as autistic individuals are all distinct people with varying interests and priorities. The popular saying within the autism community that ‘if you’ve met one person with autism then you’ve met one person with autism’ is unproblematically true on the Priority Monotropism account, as autistic individuals may have different interests and past experiences that guide the contents of their priorities.

These potential differences in the contents of the priority structures of each autistic individual may explain why there are differences in behavioural profiles between individuals. For example, it has recently been suggested that autistic girls are more likely to have special interests which are more ‘socially acceptable’ (National Autistic Society, n.d. j). This leads to them exhibiting a behavioural profile that is more difficult to identify through behavioural observation. The behavioural profiles of individuals with specific special interests will vary based on these interests.

For example, for some autistic individuals social interests may be prioritised, and for others social interests may not be prioritised. Those who prioritise social interests may not display certain social behaviours that are commonly associated with autism. These individuals may be better at identifying emotional states through facial expressions or at holding conversations with others. Social communicative differences are not wholly indicative of autism; autistic individuals can be highly sociable and may enjoy engaging in social interaction but may also struggle to understand social rules. This can be seen in the following example from the Organization for Autism Research (2016) in which the case of an autistic girl is discussed:

[Her] struggles are understanding appropriate social conduct and behaviour and respecting the personal space of others. [...] she doesn’t understand that you can’t just rock up to a stranger and sit in their lap because you think they look nice. She also doesn’t understand that sure, that man over there may be incredibly overweight but it’s probably not the nicest thing for him to hear you commenting on how he may break the chair soon if he doesn’t get up. (Organization for Autism Research, 2016)

Additionally, the topic of special interests can affect whether or not the individual is deemed to be more or less sociable by others. For example, a special interest in popular topics (such as an interest in football) may lead to fewer difficulties with fitting in than niche special interests (such as an interest in batteries).

However, it is important to note that there may still be some central common behaviours that are more likely to occur in autism as a result of narrow prioritisation in general—regardless of content—such as sensory sensitivities and difficulties with joint attention, etc.

This explanation addresses the problem of individual variance by providing the following answers to the questions that were proposed at the beginning of this subsection: why is there

individual variance in autism? Because autistic individuals are distinct people with distinct priorities that affect what they attend to – different behaviours result from these differences in priorities and attentional content. Why should we consider all autistic individuals to be members of the same group despite their potential differences in symptom profiles? Because there is a core essence, a default preference for monotropic processing, which connects them despite these differences in symptom profiles. They share a common tendency to form similar shapes of priority structures (narrow, spiky priority structures) even though there may be differences in the contents (the priorities) that are contained within those structures.

Priority Monotropism can also address an additional, hypothetical, aspect of individual variance – variance in intensity without variance in content. It is possible to conceive of two autistic individuals who have identical symptom profiles and who have active priorities that are identical in content (even though this is fairly unlikely as the contents of priorities are determined by personal interests and past experiences) and yet these individuals may still vary from each other in terms of the intensity of their symptoms.

Priority Monotropism can account for individual variance in the intensity of symptoms and behaviours, even in such a case where symptom profiles and active priorities are identical. These differences in the intensity of symptoms can be explained by differences in the frequency and intensity of monotropic processing. On the Priority Monotropism account, the intensity of monotropic processing is reflected by the shape of the priority structure – the more intense monotropic processing is, the narrower the priority structure is. Additionally, the frequency with which monotropic processing is utilised may also reflect the intensity of the symptoms and behaviours. More frequent utilisation of monotropic processing may lead to increased intensity and frequency of symptoms and behaviours.

For example, more frequent and/or more intense utilisation of monotropic processing may lead to greater levels of hypo-sensitivity (as even fewer aspects are attended to) and hyper-sensitivity (as attention to the remaining narrowly attended to stimuli is heightened). It may also create an even stronger self-reinforcing guidance loop, making task disengagement and executive functioning even more challenging; this may further reinforce engagement with RRBIs.

Section 8.4 – Accounting for Aspects That Have Been Identified by Other Essentialist Theories

As I have previously shown in chapter 2, extant essentialist theories of autism have failed to overcome the heterogeneity problem. This is something that has led to eliminativism about autism.

These essentialist theories of autism have only provided an explanation that can account for certain aspects of autism (such as sensory aspects or social aspects, etc). They cannot account for all of the aspects of autism, and so they cannot provide a sufficient answer to the metaphysical question: ‘what is autism?’.

In this section, I will demonstrate that Priority Monotropism can account for many of the aspects of autism that are central to these other essentialist theories. This will demonstrate that Priority Monotropism has the explanatory power of these other essentialist theories combined, as well as its own explanatory power (which I have already shown in this chapter and will continue to show in chapter 9).

Section 8.4.1 – Theory of Mind

As I have previously discussed in chapter 2, Theory of Mind (ToM) identifies the essence of autism as deficits in mentalising abilities. Considering this essence allows ToM to explain many of the social aspects of autism. However, as I have shown in section 8.3.1.2 of this chapter, Priority Monotropism can provide an alternative explanation for these social aspects of autism by considering differences in attentional style as the core aspect of autism.

Watzl (2023) claims that attention, specifically joint attention and attention to the attention of others, is relevant for the development of ‘mentalising’ abilities. He states that “The ability for joint attention [...] is essential for a mature understanding of other minds...” (Watzl, 2017, p.16) and that “This is how infants may first encounter other people as minded” (Watzl, 2023, p.2). This shows that, on the Priority Structure account of attention (which provides the attentional foundation for Priority Monotropism), attentional differences can explain differences in mentalising ability. This allows Priority Monotropism to account for differences in mentalising abilities in autism.

Additionally, as I have shown in chapter 2, studies that attempt to measure mentalising abilities in autism often report inconsistent findings (Gernsbacher and Yergeau, 2019).

Autistic children tend to pass modified versions of tasks that should still measure ‘mentalising’ abilities. Priority Monotropism can account for these inconsistent results, as false belief tasks often involve many different elements that must be attended to simultaneously in order to pass the test; this can result in failure to appropriately engage with the task. This explains why autistic participants are more likely to pass modified versions of these tasks (as I have shown in chapter 2).

Section 8.4.2 – Weak Central Coherence and Enhanced Perceptual Functioning

In section 8.3.1.3 of this chapter, I have shown that Priority Monotropism accounts for sensory and perceptual differences in autism, including: hyper-sensitivity, hypo-sensitivity, differences in the utilisation of multisensory processing, and differences in awareness of interoceptive sensations such as pain or hunger.

As I have previously discussed in chapter 2, Enhanced Perceptual Functioning (EPF) theory claims that the core aspect of autism is heightened perceptual functioning, which leads to sensory sensitivities in autism. Priority Monotropism can account for the central aspect of EPF by providing an alternative explanation of hyper-sensitivities. On the Priority Monotropism account, the utilisation of monotropic processing can involve the narrow prioritisation of certain sensory aspects; as a result, the individual will exhibit hyper-sensitivity to these aspects.

Additionally, Priority Monotropism does not face the issues with explaining hypo-sensitivities that EPF does. This is because, on the Priority Monotropism account, monotropic processing also involves the deprioritisation of sensory aspects that are not actively attended to. This leads to hypo-sensitivity to these aspects.

This also explains how Priority Monotropism accounts for the core aspect of Weak Central Coherence (WCC) theory. As previously discussed in chapter 2, the core aspect of WCC is that there is a local-over-global processing preference in autism (which was originally explained as resulting from a global processing deficit in autism). On this account, there is a trade-off between local and global processing abilities; local processing abilities are heightened in autism at the expense of global processing abilities. This leads to hyper-sensitivity to local sensory aspects and hypo-sensitivity to global sensory aspects.

As I have already discussed in chapter 2, WCC is evidenced by local and global processing tasks such as Navon tasks, embedded figures tasks, and visual search tasks. These tasks measure the individual's ability to focus on local aspects of a scene by disregarding the global aspects of that scene. Priority Monotropism can provide an alternative explanation that allows it to account for the results of these tasks while also accounting for inconsistencies that have been reported in the findings of these tasks (as I will explain later in chapter 9). Additionally, Priority Monotropism can do so without having to rely on the terms 'local' and 'global' (which I have argued are potentially problematic in chapter 2).

Section 8.4.3 – Executive Dysfunction

In section 8.3.1.1 of this chapter, I have provided an explanation of how executive dysfunction occurs in autism and the role that it plays in explaining RRBI. The guidance processes that are involved in attention on the Priority Structure account are utilised by the Priority Monotropism account to explain how autistic individuals may often find themselves stuck in a cyclical guidance loop. This causes the sticky attention that is commonly reported in autism. This sticky attention makes it difficult for the individual to disengage from their current task or priority in order to switch to a new task or priority. This allows Priority Monotropism to account for the essence identified by Executive Dysfunction theory while also being able to account for a wide range of symptoms and behaviours that this theory was not able to account for (such as perceptual differences, which I have discussed in chapter 2).

Section 8.5 – Chapter Conclusion

In this chapter, I have introduced a new theory of autism, Priority Monotropism, which combines central elements from the extant Monotropism theory of autism and the Priority Structure account of attention.

I have shown that Priority Monotropism is able to address the heterogeneity problem of autism by accounting for a wide range of symptoms and behaviours that are common to autism and by accounting for individual variance in symptom profiles. This has allowed Priority Monotropism to push back against autism eliminativism by showing that autism can be meaningfully understood as a unified condition.

Additionally, in section 8.4, I have explained that Priority Monotropism can account for the central aspects of autism that are identified by the essentialist theories that I have previously outlined in chapter 2. In doing so, I have demonstrated that Priority Monotropism has the explanatory power of these essentialist theories combined, as well as its own additional explanatory power.

In the next and final chapter, I will explore potential future directions for research relating to this theory. In doing so, I will further consider the potential explanatory power of this account by briefly considering how Priority Monotropism could potentially account for other aspects of autism and autism research that have been difficult to account for. I will consider why it is the case that autism frequently co-occurs with a wide range of other conditions, why there have been inconsistencies in autism research findings, and finally, how this theory has laid an improved foundation for future empirical research efforts and support efforts.

Chapter 9 – Directions for Future Research

Section 9.1 – Chapter Introduction

In the last chapter, I outlined a novel essentialist theory of autism: Priority Monotropism. I have shown how this theory combines elements from the extant account of Monotropism and the Priority Structure account of attention in order to create a revised and improved version of Monotropism. I demonstrated the improved explanatory power of this account by showing that it is better equipped to resolve the heterogeneity problem than the other essentialist theories of autism (which were previously discussed in chapter 2).

In this chapter, I will consider future directions for research. I will briefly consider how Priority Monotropism may be able to explain many other aspects of autism that remain unexplained and have been difficult to account for. In section 9.2, I will provide some ideas about how Priority Monotropism could potentially explain why autism co-occurs with a wide range of other conditions. In section 9.3, I will examine how Priority Monotropism could potentially explain inconsistencies in autism research findings. Finally, in section 9.4, I will claim that Priority Monotropism provides an improved foundation for future investigative efforts into empirical testing and practical support.

However, it is important to note that the purpose of this chapter is to lay out these ideas for future research. As such, I will not be providing complete explanations of how Priority Monotropism addresses these issues. Further research into these complex issues would be required to make any definitive claims, and such research efforts are beyond the scope of this thesis. Instead, I aim to further demonstrate the potential explanatory power of this account by providing suggestions about how future research into autism through the lens of Priority Monotropism could help to illuminate these issues.

Section 9.2 – Accounting for Co-Occurring Conditions

As I have previously discussed in chapter 1, autism often co-occurs with other conditions. These co-occurring conditions are conditions that may often occur alongside autism, but which are distinct, separable conditions that may also occur without autism. Chung and Kim (2024) report that 95% of autistic individuals also have at least one other co-occurring condition.

As I will demonstrate in this chapter, there are a wide variety of conditions that commonly co-occur with autism. Some of these co-occurrences are fairly simple to account for (such as anxiety and stress, which I will discuss in section 9.2.8), while others are more difficult to explain (such as synaesthesia or agnosia, which I will discuss in sections 9.2.6 and 9.2.4, respectively). In this section, I will briefly provide ideas for how Priority Monotropism could account for the co-occurrence between autism and the following conditions: savant syndrome, hyperlexia, dyspraxia, agnosia, alexithymia, synaesthesia (including mirror-touch synaesthesia), stress and anxiety disorders, and somatic conditions. Then, in section 9.2.10, I will discuss the explanatory limitations of Priority Monotropism. I will discuss the fact that the co-occurrence of some conditions cannot be accounted for at the level of explanation that I am proposing in the development of the Priority Monotropism account.

Section 9.2.1 – Savant Syndrome

Savant syndrome is a condition in which an individual has an extremely uneven profile of skills. Individuals with savant syndrome are characterised by their demonstration of extreme skill levels in a highly specific area while also having a considerable intellectual disability. Treffert describes the condition as follows: “Savant syndrome is a rare, but extraordinary, condition in which persons with serious mental disabilities, including autistic disorder, have some ‘island of genius’ which stands in marked, incongruous contrast to overall handicap.” (Treffert, 2009, p.1351). He claims that 1 in 10 autistic individuals display savant skills (Treffert, 2009).

However, it is important to note that savant syndrome is a separate condition from autism that can also affect non-autistic individuals. While savant syndrome commonly co-occurs with autism, not all individuals with savant syndrome are autistic (Treffert, 2009).

Howlin et al. (2009) provide the following list of commonly reported savant skills:

...mathematical skills (calendrical calculations, lightning arithmetic and prime number calculations), music (especially the ability to replay complex sequences after only one exposure), art (complex scenes with accurate perspective either created or replicated following a single brief viewing) and memory for dates, places, routes or facts [...] ‘pseudo-verbal’ skills (hyperlexia or facility with foreign languages), coordination skills and mechanical aptitude. (Howlin et al., 2009, p.1360)

Hyperlexia is described in this list as a ‘pseudo-verbal’ skill as it relates to the ability to read written language despite difficulties with verbal communication and interpreting meaning. I will discuss this further in the next section, which is dedicated to the discussion of hyperlexia.

Some people have viewed the existence of savant syndrome as raising a paradox – they have questioned how it is possible for someone with considerable intellectual deficits to display talents that go well beyond the norm.

However, considering these savant skills through the lens of Priority Monotropism shows that they are not paradoxical. On this account, savant syndrome could be framed as an extremely spiky profile, in which priority structures may be even narrower and/or monotropic processing may be more frequently utilised. This notion of an extremely spiky profile ties in well with the core aspects of savant syndrome (extreme skills in one or a few areas at the expense of skills in a wide range of areas). On this view, the area of skill is even more highly prioritised while an even greater range of potential priorities are deprioritised.

This increased utilisation of monotropic processing would even further influence the guidance principles that guide attention (which I have outlined in chapters 7 and 8). For example, the area of skill being more highly prioritised by the individual will affect its value, which in turn will promote stronger attentional guidance to that area of skill, while the exclusionary nature of monotropic processing will influence the guidance principle of ‘selection history’ and ‘general world knowledge’ (since, as fewer things are attended to, there is less variety in these areas). This leads to a cyclical guidance loop, which is responsible for aspects of RRBIs in autism such as insistence on sameness and sticky attention (as I have previously discussed in chapter 8). If savant skills are even more highly prioritised than areas of special interest typically are in autism, then the guidance loop would be even tighter and more restricting. This could lead to increased repetition and obsession with the area of skill—leading to increased levels of practice and engagement with said skill.

Additionally, restrictions in the range of available interests, due to co-occurring intellectual disability, may also further restrict the range of interests that are available to individuals with savant syndrome. This further narrows the range of accessible potential priorities from the start, before the cyclical guidance loop is even initiated. For example, Treffert (2009) discusses the case of Thomas Fuller, a savant with incredible skill in calculation. It was claimed that he:

... ‘could comprehend scarcely anything, either theoretical or practical, more complex than counting’. However, when Fuller was asked how many seconds a man had lived who was 70 years, 17 days and 12 hours old, he gave the correct answer of 2210500800 in 90s, even correcting for the 17 leap years included. (Treffert, 2009, p.1351)

In this example, I am proposing that Fuller’s intellectual disability restricted his range of potential priorities, as there were serious limitations in what he would have been able to effectively engage with.

This ties in well with evidence that suggests that savant skills are developed or strengthened through obsessive practice and engagement.

“Intense interests can lead to a quasi-exclusive dedication of time and cognitive resources to the material of interest that facilitates skill acquisition in specific domains, thus contributing to the emergence of special abilities” (Ostrolenk, Boisvert, and Mottron, 2025, p.8)

A study by Hughes et al. (2018) reported that autistic savants displayed higher rates of obsessional traits when compared to autistic non-savants. They claim that this supports the theory, which was proposed by Simner, Mayo, and Spiller (2009), that savants have an “obsessive urge to over-rehearse their skills to prodigious levels.” (Hughes et al., 2018, p.2).

Additionally, Hughes et al. (2018) conducted a study in which autistic savants and autistic non-savants were both tasked with performing a novel common savant skill: calendar calculation. They reported that the autistic savants did not immediately display an advantage on this task over the non-savant autistic participants. This reinforces the idea that obsessive practice plays an important role in the development of savant skills. Additionally, they reported that the autistic savants appeared to display a more engaged and effortful approach to the task – showing that they engaged in the task with greater intensity. This approach of greater intensity may relate to the development of savant skills; the fact that no such skills were developed despite the utilisation of this approach may be due to the fact that more time was needed to practice. On this view, such skills develop from obsessive practice, which requires time.

Additionally, a study by Casey et al. (1993) reported attentional differences, which would be predicted by Priority Monotropism, in autistic savants. They claim that autistic savants were

significantly less effective at dividing attention, especially for the utilisation of multimodal processing, than the control subjects, stating that:

A striking inability to efficiently divide attention between modalities was observed during simultaneous performance of a visual and auditory detection task. The [autistic] savants failed to detect targets significantly more than did the controls. The [autistic] savants appeared to focus on one modality at the expense of failing to detect targets presented in the other modality. The ignored modality appeared to vary across subjects and changed within some subjects over trials, indicating that, if given enough time, the [autistic] savants could eventually shift their attention to the relevant modality. (Casey et al.,1993, p.942)

They also reported the autistic savants as displaying: “a deficit in the ability to disengage attention from one spatial location to another...” (Casey et al.,1993, p.943) and stated that:

In other words, when in an alert state, but when attention is not yet engaged, any target of sufficient intensity has little problem capturing attention. However, once the attentional system is engaged, detection of another stimulus requires that that stimulus have enough intensity or significance to draw attention from the original stimulus to a new one. With savants, once attention is engaged, available resources for detecting other signals or stimuli may be depleted or inhibited. This inhibition or refractoriness may continue for long intervals of time. (Casey et al.,1993, p.944)

This ties in well with my explanation about guidance and cyclical guidance loops that lead to ‘sticky attention’ (which was previously discussed in chapter 8). Attentional guidance will be restricted and repetitive until something with higher guidance values (something which is more salient or more valued) interrupts the guidance loop.

Casey et al. (1993) also note the importance of repetition for the development of savant skills through their discussion of the common savant skill, calendar calculation. They state that repetitive engagement with the skill can lead to mastery and that skills that have repetitive natures are more likely to be attractive to autistic individuals who often display an insistence on sameness:

The calendar has a very repetitive nature (e.g., in days of the week and months in a year). In fact, the calendar repeats itself every 28 years. One may better understand how overselectivity or overfocusing on a task of such repetition could result in its ultimate mastery even for moderately functioning individuals. The repetition observed in the calendar may also explain its attractiveness to autistic individuals who have been observed to insist on order and sameness. (Casey et al., 1993, p.944)

However, it is important to note the limitations of this study by Casey et al. (1993). In this study, autistic savants were compared to non-autistic non-savant controls. While these findings support Priority Monotropism in general, and also importantly suggest that monotropic processing is utilised by autistic savants as well as non-savant autistic individuals, this study does not tell us about whether autistic savants have heightened attentional differences when compared to non-savant autistic individuals. Conducting such a study (which uses non-savant autistic individuals as the control group) in future would be beneficial for further examination of this issue.

It is important to note that this explanation of the development of skills in autism through the utilisation of monotropic processing is not exclusive to cases of savant syndrome. Priority Monotropism also has the potential to explain cases of increased skill development in autism in general (without savant syndrome).

As I have previously discussed in chapter 8, on the Priority Monotropism account, autistic individuals have a ‘spikier’ priority profile due to monotropic processing. This allows for a similar explanation for the development of skills to that which I have provided for savant skills.

This explains why autistic individuals develop skills and expertise in relation to their areas of special interest. These special interests are characterised by:

- (a) significant depth and breadth of knowledge about the area, (b) difficulty attending to other important information in the environment due to preoccupation with the area, (c) engagement in prolonged periods of study related to the area, and (d) a strong intensity in focus uncommon for neurotypical peers. (Bross and Travers, 2017, p.75)

Woods and Estes (2023) specifically link the development of skills relating to special interests to monotropic processing: “Because autistic people tend to experience monotropism,

they may be able to devote long periods of time to studying or talking about one subject, which often leads to expertise and mastery.” (Woods and Estes, 2023, p.2)

However, it is important to note that the development of *specific* savant skills is not something that is explained solely by Priority Monotropism. As I have previously discussed in chapter 8, autistic individuals have distinct priorities. Priority Monotropism does not predict the *specific* contents within the individual's priority structure (it does not predict what they will specifically prioritise) but rather how they will prioritise said content (it predicts how attention will be distributed to that content). What makes someone specifically interested in developing certain skills or pursuing specific interests (such as music or art) may depend on other factors that determine their interest or aptitude (such as having perfect pitch or having photographic memory). These are factors that often overlap with savant skills, which may make that skill a good fit for the individual. However, as I have shown, Priority Monotropism explains increased, obsessive and repetitive engagement with these skills at the expense of other potential interests; this is something that may further motivate and enable the development of these skills.

Additionally, it is important to note that although Priority Monotropism explains how spiky profiles may lead to the development of savant skills and explains how these spiky profiles might be cyclically reinforced, it does not explain why some individuals have spikier profiles than others in the first place.

Section 9.2.2 - Hyperlexia

Ostrolenk et al. (2017) define hyperlexia as:

... the co-occurrence of advanced reading skills relative to comprehension skills or general intelligence, the early acquisition of reading skills without explicit teaching, and a strong orientation toward written material, generally in the context of a neurodevelopmental disorder. (Ostrolenk et al., 2017, p.134)

Under this definition, hyperlexia may be understood as a specific kind of savant skill.

It is a condition that is characterised by “an intense fascination with letters or numbers or, in younger people, an ability to read far beyond their age.” (National Autistic Society, n.d. k)

Hyperlexia is a distinct condition that is separable to but commonly co-occurs with autism. Around 6%-14% of autistic children have hyperlexia (WebMD, 2025) and around 84% of children with hyperlexia are autistic (Ostrolenk et al., 2017). This shows a significant overlap between the two conditions.

Autistic children with hyperlexia are described as being "...obsessed with numbers and letters, preferring books and magnetic letters over other types of toys. They also frequently remember important numbers such as license plates and birth dates." (WebMD, 2025)

However, despite these improved written language skills, autistic individuals with hyperlexia often have difficulty with communicating and understanding verbal language (National Autistic Society, n.d. k). This shows that hyperlexia is not necessarily linked to improved language skills in general.

Priority Monotropism would explain the co-occurrence of autism and hyperlexia through the same explanation that was given in the previous section for savant skills in general. The utilisation of monotropic processing can lead to a cyclical guidance loop that reinforces a restricted range of interests. The individual will spend more time with these interests and may develop superior skills relating to them.

This idea, that hyperlexia is reinforced by RRBIs and increased time dedicated to it, is expressed by Solazzo et al. (2021), who state that:

The early ability to name and recognize letters and numbers was associated with a higher level of restricted and repetitive behavior [...] Moreover, it has been suggested that in children with ASD and hyperlexia, a strong interest in reading results in an increase in time and practice devoted to this ability, thus leading to the emergence of advanced reading skills. [...] As for many RRBs, this early interest in literacy could be considered invasive as it may prevent children from engaging in other activities... Solazzo et al. (2021, p.1-2)

Additionally, Priority Monotropism may help to explain why hyperlexia is one of the most common savant skills in autistic children. Interestingly, hyperlexia presents a case of early enhanced skill acquisition, with the earliest cases typically emerging at around 2 years old (Solazzo et al., 2021; Ostrolenk, Boisverta, and Mottron, 2025). Many young children may have access to letters and numbers through educational toys and children's books as well as

objects in their general environment. This could allow for this to be a common early interest on the Priority Monotropism account.

As I have previously discussed in chapter 8, young autistic children may not have access to a wide range of potential interests until they become less interdependent and develop enough to gain the skills required to engage with further potential interests. This could potentially explain why many autistic children with hyperlexia abandon their interest in written words, letters, and numbers in order to pursue other interests at around age 4, when they may start to gain access to a wider range of potential interests to pursue (Ostrolenk, Boisverta, and Mottron, 2025).

As I have previously discussed in chapter 8, Priority Monotropism accounts for the fact that special interests often rotate in autism and that special interests can become exhausted, which allows for a new interest to take centre stage. Priority Monotropism could explain this abandonment of interests related to hyperlexia by stating that written language may become an exhausted interest due to over-engagement. This idea is supported by Ostrolenk, Boisverta, and Mottron (2025), who state that:

The interest often peaked around age 2, and the reduction appeared around age 4, drastically in more than one case. The reduction was also generally associated with the development of new interests, suggesting that the interest may have expanded to other themes once the child had fully explored it, rather than disappeared. (Ostrolenk, Boisverta, and Mottron, 2025, p.5)

Section 9.2.3 - Dyspraxia

As I have previously shown in chapter 2, autistic individuals often experience difficulties with co-ordination and integration. In chapter 8, I have explained that Priority Monotropism accounts for these difficulties through the utilisation of monotropic processing, which makes co-ordination and integration more difficult for autistic individuals. Monotropic processing makes it more difficult to successfully attend to the many different sensory cues that aid balance and co-ordination abilities (for example, to attend to relevant visual cues and proprioception simultaneously).

Co-ordination and integration are essential elements of the skills that are impaired in dyspraxia (such as hand-eye co-ordination, difficulties with motor functioning – fine and

gross, and balance) (Dyspraxia UK, n.d.). Dyspraxia is significantly more common in autistic populations than in non-autistic populations, affecting around 7% of autistic individuals and only around 1% of non-autistic individuals (Cassidy et al., 2016).

The co-occurrence of autism and dyspraxia is potentially well explained by Priority Monotropism, as monotropic processing can lead to these co-ordination and integration difficulties that characterise dyspraxia.

Section 9.2.4 – Agnosia

Agnosia is a separate condition, which may co-occur with autism, that is broadly defined as “an inability to recognise the input of sensory information.” (Isaacs, 2026). This can lead to difficulties with identification.

There are different types of visual agnosia, including:

- Object agnosia – difficulties with identifying objects
- Topographic agnosia – difficulties with identifying locations
- Prosopagnosia – difficulties with identifying faces

(Behrmann, 2026)

Priority Monotropism could potentially explain agnosia as resulting from monotropic processing. As I have already discussed in chapters 3 and 8, narrow attentional distribution can make it more difficult to detect and integrate all of the perceptual cues that are needed to form relevant associations.

Bogdashina (2016) supports this idea with the example of Paul Isaacs’s (an autistic individual and autism advocate) experiences of prosopagnosia:

“...due to visual fragmentation, he cannot process faces as wholes, and the only reference he has for people is their hair ‘because that makes up the largest coherent whole that isn’t as difficult as the face’” (Bogdashina, 2016, p.162)

Additionally, a broader case of visual agnosia has already been explained by Monotropism (in chapter 3) through the discussion of Temple Grandin and her experiences of difficulties with identifying animals (relying solely on single features instead of gestalts to identify animals).

Section 9.2.5 - Alexithymia

Alexithymia, also known as ‘affective agnosia’ (Lane et al., 2020), is a condition that is “characterized by an impaired ability to be aware of, explicitly identify, and describe one’s feelings.” (Hogeveen and Grafman, 2021, p.48).

A meta-analysis by Kinnaird, Stewart, and Tchanturia (2019) estimated that the rate of co-occurrence between autism and alexithymia is between 33% and 63%.

Priority Monotropism provides a potential explanation for this co-occurrence between autism and alexithymia. In chapter 8, I provided an explanation for how Priority Monotropism could account for interoceptive differences in autism. Hyper-awareness of interoceptive states is caused by narrow attendance to these states, while hypo-awareness of interoceptive states is caused by narrow attendance to external states. Priority Monotropism explains alexithymia as a further potential consequence of these interoceptive differences. Widespread lack of awareness, which results from monotropic processing, leads to difficulties with identifying one’s own emotions, just as it leads to difficulties with identifying other interoceptive states.

This idea is supported by the ‘attention-appraisal’ model of alexithymia in which attentional differences, including reduced reliance on interoceptive attention when attention is externally directed, lead to difficulties with the identification of one’s own emotions (Preece et al., 2017). Additionally, this is also supported by a claim made by Kinnaird, Stewart, and Tchanturia (2019, p.87) that autistic individuals “direct fewer attentional resources towards emotional stimuli”.

Section 9.2.6 - Synaesthesia

Synaesthesia is a condition in which typical sensory experiences elicit the automatic pairing of an additional ‘illusory’ (synaesthetic) experience (Banissy et al., 2013). Synaesthesia comes in many different forms and can be intramodal or cross-modal in nature. Colour-grapheme synaesthesia (in which experiences of words, letters, and numbers elicit additional synaesthetic experiences of colours) provides an example of intramodal synaesthesia. Audio-

visual synaesthesia (in which experiences of sounds elicit additional synaesthetic visual experiences) provides an example of cross-modal synaesthesia.

Synaesthesia is significantly more prevalent in autism populations. It affects around 4% of the general population (Simner et al., 2006) and around 18% of the autistic population (Baron-Cohen et al., 2013).

In my Master's dissertation (Annan, 2021), I outlined a novel explanatory account of why synaesthesia may occur more commonly in autism populations than in non-autism populations. I was particularly motivated to explain this co-occurrence as core aspects of autism, such as reduced perceptual integration (both intramodal and multisensory), may initially seem to be in direct opposition to synaesthesia, which is characterised by increased perceptual integration (both intramodal and multisensory).

The account that I developed was titled the 'Monotropic-Compensatory' account. I explained that, on this account, (multisensory and/or intramodal) synaesthesia may develop in order to compensate for the underutilisation of multisensory and/or global processing in autism, which occurs as a result of monotropic processing.

I supported this account by providing evidence that traits relating to monotropic processing were linked to synaesthesia. For example, Ward et al. (2017) reported that non-autistic synaesthetes displayed elevated Autism Quotient (AQ) scores. The AQ is a diagnostic self-report questionnaire that is used to measure autism traits. The following five traits are assessed by the AQ: social skills, attention switching, attention to detail, communication, and imagination. However, Ward et al. (2017) found that only the 'attention to detail' subscale of the AQ was related to these increased scores, suggesting that increased attention to detail may be a potential link between autism and synaesthesia. Similar findings were also reported by van Leeuwen et al. (2019). Additionally, Ward et al. (2017) reported increased rates of both hyper- and hypo- sensory sensitivities in non-autistic synaesthetes. I then presented evidence that demonstrates that there is reduced multisensory integration in synaesthesia by explaining that synaesthetes have shown reduced susceptibility to multisensory illusions, such as the sound induced-flash illusion (Neufeld et al., 2012) and the McGurk effect (Sinke et al., 2012).

I then provided evidence that demonstrates that it is possible for compensatory synaesthesia to occur. I did this by showing how unimodal sensory deficits can lead to the development of synaesthesia in the affected modality (for example, the development of visual synaesthetic

experiences in response to visual deficits). For example, Steven and Blakemore (2004) found that late blind participants developed compensatory forms of visual synaesthesia (including auditory-visual chromesthesia, in which synaesthetic experiences of colour were induced in response to sounds, and tactile-visual synaesthesia, in which the tactile reading of braille induced visual experiences of colour). I also provided evidence to support the idea that it is possible for multisensory deficits to elicit similar compensations. Stevenson et al. (2017) have shown that reduced global/multisensory integration affects the intensity of hallucinations in schizophrenia. Additionally, Williams et al. (2010) found that schizophrenic participants who experience both audio and visual hallucinations displayed greater deficits in multisensory processing than schizophrenic participants who experience only auditory hallucinations.

I took this evidence together to suggest that it is possible that synaesthesia may result more commonly in autism populations as the result of a compensatory process that compensates for the reduced utilisation of ‘global’ and multisensory processing, which results from monotropic processing.

I aim to continue to examine and develop this account in future by further considering it through the lens of Priority Monotropism.

Section 9.2.7 – Mirror Touch Synaesthesia

Mirror-touch synaesthesia is another form of synaesthesia that can co-occur with autism (Baron-Cohen et al., 2016). I have not considered mirror-touch synaesthesia in the account that I developed throughout my Master’s dissertation.

Ward and Banissy (2015) describe mirror-touch synaesthesia as “... the conscious experience of tactile sensations induced by seeing someone else touched.” (Ward and Banissy, 2015, p.118)

Priority Monotropism could also potentially explain mirror-touch synaesthesia through the ‘Monotropic-Compensatory’ account. As I have previously mentioned in chapter 8, monotropic processing typically leads to reduced awareness of one’s own interoceptive states (such as pain, hunger, and emotions). I have shown in chapter 8 that autistic individuals often externally direct their attention and rely more heavily on external visual cues in order to determine their interoceptive states. For example, seeing a wound or seeing broken glass on

the ground in order to determine that they are in pain. These deficits in interoceptive awareness may plausibly be compensated for by the development of mirror-touch synaesthesia—in which the detection of visual cues relating to the interoceptive states of others may mistakenly be attributed to the observer.

On this view, autistic individuals may develop mirror-touch synaesthesia through an overreliance on external cues when determining their own interoceptive states. As I have shown in chapter 8, this overreliance can be explained by monotropic processing.

Pronoun reversal provides additional evidence that may support this idea. Autistic individuals may demonstrate pronoun reversal, in which pronouns used to identify and distinguish oneself from others (such as ‘I’ and ‘you’) are reversed (Overweg, Hartman, and Hendriks, 2018). For example, an autistic child may say “you want a snack” instead of “I want a snack” when making a request. This may suggest potential confusion or blurring between the self-other boundary, which may explain why autistic individuals with mirror-touch synaesthesia might mistakenly attribute the visual cues associated with the interoceptive states of others to themselves.

Section 9.2.8 – Stress and Anxiety

Anxiety disorders are significantly more common in autism populations than in non-autism populations. Srinivasan (2025) states that anxiety disorders affect 18% of the general population and 69% of the autism population.

As I have shown throughout this thesis, autistic individuals exhibit a wide range of symptoms and behaviours that can affect their ability to make and maintain social connections, perform tasks related to daily functioning, secure employment, and complete education. This is especially the case when sufficient support is not provided. As I have shown in chapter 1, these issues are likely to cause distress.

Additionally, the utilisation of masking may further increase stress. This is because, on the Priority Monotropism account, masking requires working against monotropic attentional guidance. When autistic individuals mask, they may need to adopt a more polytropic processing style in order to attend to as many relevant social cues as possible. As previously discussed in chapter 3, this change in attentional style requires additional effort. Furthermore, masking may require the individual to fight against the strong attentional guidance relating to

their interests. This is to say that masking may require the individual to utilise executive control in order to forcibly change the target of their attention.

Srinivasan (2025) states that:

In response to social pressures, many autistics engage in effortful social adaptation, suppressing natural traits to conform to neurotypical expectations. This may involve mimicing social behaviors, forcing eye contact, or masking sensory sensitivities to minimize rejection. [...] Though these strategies may reduce immediate discrimination, they require significant cognitive effort, leading to exhaustion, identity erosion, and increased anxiety over time. (Srinivasan, 2025, p.5)

Furthermore, as I have previously discussed in chapter 8 through my discussion of ‘meltdowns’, monotropic processing may affect the individual’s awareness of their stress levels, making it difficult to tell that stress is building in order to effectively manage it before it becomes overwhelming.

Srinivasan (2025, p.3) states that “... altered interoception, difficulties perceiving internal bodily states, can hinder emotional regulation, making it harder to identify and manage stress.”

Conversely, attention to interoceptive states (such as one’s own heart rate) during monotropic processing may lead to increased awareness and anxiety in relation to those states. This heightened awareness may cause the individual to experience further anxiety and enter into a ‘feedback loop’ of anxiety. This is shown in the following self-reported experiences from two different autistic individuals in a study by Adams et al. (2025):

Like normally the path will be I’ll get anxious about something and then um ... I’ll go to break and it will be really noisy and that leads to me getting overwhelmed and then I’ll go to another lesson and that’s noisy and that escalates and it’s just like too much and then I become more aware of everything like my heart my breathing and then I’ll like shut down. (Adams et al., 2025, p.1590)

I’ll find that the increase in heartbeat ... definitely creates a feedback loop of anxiety. (Adams et al., 2025, p.1590)

Additionally, autistic individuals are more likely to experience ‘catastrophic’ thoughts and worries (National Autistic Society, n.d. 1) that contribute to stress and anxiety. Priority Monotropism could potentially explain this as occurring due to attention getting stuck on these negative thoughts during monotropic processing. This idea is supported by the following explanation from the National Autistic Society (n.d. 1): “Catastrophising is when you are anxious about a situation and that anxiety leads your brain to ruminate on a potentially disastrous outcome.” This aspect of rumination is well explained by Priority Monotropism, as sticky attention may make it difficult to stop attending to anxious thoughts and feelings.

Section 9.2.9 – Somatic Conditions

Somatic conditions are characterised by the having of “...one or more physical symptoms accompanied by an excessive amount of time, energy, emotion, and/or behavior related to the symptom that results in significant distress and/or dysfunction.” (D'Souza, and Hooten, 2023, p.1). These conditions are significantly more prevalent in autism populations (Williams, and Gotham, 2022).

Importantly, a study by Carpita et al. (2024) found that non-autistic participants with a somatic condition displayed elevated AQ scores, but only within the ‘attention switching’ domain. They claim that:

... patients suffering from [somatic conditions] were reported to show significantly higher attention switching difficulties compared to healthy controls, which may be explained by cognitive inflexibility that may ultimately lead to fixation on bodily sensations, heightening the intensity of somatic symptoms. Secondly, it has been hypothesized that alterations in sensitivity to sensory stimuli can possibly influence the perception of somatic symptoms. Indeed, autistic people’s interoceptive systems are known to function differently, leading them to exhibit both an increased sensitivity to physiological sensations and a decreased capacity to accurately detect bodily signals. Based on these premises, the higher interoceptive sensibility of people with ASD may contribute to the greater prevalence of somatic symptoms in said population. (Carpita et al., 2024, p.6-7)

Additionally, another study by Huang et al. (2025) found that only the attention to detail subscale was responsible for greater need for reassurance in relation to somatic experiences in autistic individuals. They claim that:

These findings extend earlier knowledge of the correlation between somatic symptoms and overall autistic traits with the more specific domain of attention allocation. Autistic individuals who have difficulties with attention allocation may tend to focus on trivial changes of bodily sensation; such constant over-concerns often require more explanations and psychological support to relieve the anxiety. (Huang et al., 2025, p.7)

Priority Monotropism is well suited to explain this increased prevalence of somatic conditions in autism. On this account, the cyclical guidance of attention during monotropic processing may explain why autistic individuals are more likely to experience somatic conditions. This cyclical guidance of attention leads to sticky attention and difficulty with switching attention in order to disengage from the pain/distress caused by the physical symptom(s). Additionally, as I have previously discussed in chapter 8 and earlier in this chapter, when attention is internally directed during monotropic processing, interoceptive experiences may be more intense and anxiety-inducing; this may lead to increased intensity of somatic experiences and further anxiety-driven fixation.

Section 9.2.10 – Explanatory Limitations

However, while Priority Monotropism is well equipped to potentially explain many co-occurring conditions, it is not well equipped to explain all of them. For example, co-occurring medical conditions, such as epilepsy and gastro-intestinal issues, are not explained by Priority Monotropism. These conditions may be related to a neuronal or biological marker of autism, which may be identified in the future.

Besag (2017, p.1) claims that “There is no plausible mechanism for autism causing epilepsy, however. The co-occurrence of autism and epilepsy is almost certainly the result of underlying factors predisposing to both conditions, including both genetic and environmental factors.”

The fact that the explanation for these co-occurrences could lie at the neurological or biological level is unproblematic for Priority Monotropism, as this account does not make any claims about such markers.

Priority Monotropism can account for heightened levels of anxiety and stress, which could explain increased difficulties with gastro-intestinal issues. However, it does not necessarily account for the instantiation of such issues.

Additionally, as I have shown in the previous section, Priority Monotropism may explain increased rates of somatic conditions due to heightened interoception. This is also something that could relate to increased perceived intensity of gastro-intestinal issues.

Section 9.3 – Explaining Inconsistencies in Autism Research

In this section, I will briefly consider the ways that Priority Monotropism could potentially resolve inconsistencies in autism research. I will explain that Priority Monotropism can provide a unifying explanation for the following seemingly opposing views within autism research: inconsistencies concerning local/global processing abilities and inconsistencies concerning expectations and priors.

Section 9.3.1 – Inconsistencies in Local/Global Task Performance

Differences in local/global task performance have been suspected to be an important aspect of autism (as I have shown in my discussion of Weak Central Coherence theory in chapter 2).

However, Koldewyn et al. (2013) reported that the findings from empirical studies on these abilities are conflicting. In order to further examine these abilities, they conducted a study in which autistic children were tested on their abilities to identify local/global levels of a Navon task image (see figure 2.2 in chapter 2 for an example). They utilised two task conditions, one in which the participants were explicitly instructed to attend to a specific level (local or global) and one in which the participants were not instructed on how to attend to the task. The results of this study found that participants were able to attend to the global-level stimuli when they were explicitly instructed to do so, but that they would typically attend to the local-level stimuli when they were not instructed on how to attend.

These findings are consistent with, and well explained by, Priority Monotropism, as on this account, the default utilisation of monotropic processing (in which global-level stimuli are

more likely to be disregarded) stems from a preference as opposed to a deficit in global processing ability.

Section 9.3.2 – Inconsistencies in Explanations of Expectations and Priors

As I have previously discussed in chapter 2, one of the common symptoms of autism is insistence on sameness/aversion to change. One proposed explanation for these symptoms is that autistic individuals form expectations differently, which may lead to many things feeling unexpected or shocking. As such, autistic individuals may prefer to follow strict routines and avoid changes in their environment in order to combat this uncertainty by increasing predictability. However, competing explanations of how these differences in expectations form have been suggested.

Pellicano and Burr (2012) have claimed that autistic individuals form hypo-priors, in which past experiences and knowledge are under-relied on. They claim that this allows autistic individuals to have a more accurate experience of the world—unchanged by their expectations. However, this may also lead to things often being surprising – as expectations are not being sufficiently formed.

However, Van de Cruys et al. (2013) propose an opposing explanation. They claim that autistic individuals form hyper-priors, in which past experiences and knowledge are rigidly over relied on. These strong and rigid expectations are often overfitted, meaning that they are applied to all future scenarios even when they are not appropriate. For example, after seeing a dog in a specific café on one occasion, the individual may always expect to see a dog in every café. On this account, this overfitting leads to surprise, as the individual had a strong expectation that was proven to be wrong.

However, it may be possible to unify these explanations by considering this issue through the lens of Priority Monotropism. On this account, it could be the case that autistic individuals will form hypo-priors in relation to things that are unattended and hyper-priors in relation to things that are narrowly and intensely attended to. This explanation would also potentially fit in with the picture of cyclical guidance that I have previously discussed in chapter 8. Autistic individuals may rely more heavily on things that have previously been attended to (such as special interests) and as a result may form greater levels of knowledge and expectations about those things. However, they may not form expectations and knowledge (such as general world knowledge) about things that are not attended to in this way.

Section 9.4 – Improved Foundations for Future Examination of Empirical Testability and Practical Support

In chapter 3, I have previously claimed that no specific tests have been developed to empirically test the extant Monotropism account. This is something that may hinder its explanatory power, as other theories of autism, such as the Weak Central Coherence theory and the Theory of Mind theory, are supported by empirical findings from tasks that are designed to measure their essence (for example, false belief tasks, which I have discussed in chapter 2). While I have argued (in chapter 2) that these tasks do not always provide an effective measure for examining these accounts, it is still a benefit of these accounts that such tests can be developed and refined. In chapter 3, I have argued that the main reason that the extant account of Monotropism was not empirically testable is because of the fact that central claims about ‘attention’ and ‘interest’ were too vague to test. The extant authors have not told us how these terms should be considered on this account, and as such, it is unclear what we should be measuring through an empirical test.

In chapter 8, I have shown that Priority Monotropism improves upon the extant account by making clearer claims about attention, interest, and the nature of monotropic processing. As such, it has provided an improved foundation for the future examination of how this account might be empirically tested. While providing any specific suggestions for such an empirical test is beyond the scope of this thesis, I aim to examine this issue in subsequent works.

Additionally, providing clearer claims may help to improve our understanding of autism and the challenges that autistic individuals may face in areas such as education and employment. However, again, identifying specific support methods is beyond the scope of this thesis.

Section 9.5 – Chapter Conclusion

In conclusion, in this chapter I have demonstrated that Priority Monotropism has further potential explanatory power over and above that which I have shown in the previous chapter. I have shown that Priority Monotropism may allow for an explanation of why autism commonly co-occurs with a wide range of, often, seemingly unrelated conditions. I have also shown how considering areas of current conflict in autism research through the lens of Priority Monotropism may allow for a unified resolution of these conflicts. Additionally, I have claimed that Priority Monotropism has improved upon the extant account by laying a

foundation for future investigation into empirical testing and improved support methods.

However, it is important to note that all of the ideas in this chapter are still at an early stage of development. I will continue to consider these issues in future works.

Conclusion

In conclusion, throughout this thesis I have developed a novel account of autism – Priority Monotropism. Priority Monotropism is an essentialist theory of autism that identifies the essence of autism as attentional differences. Much like the extant account of Monotropism, Priority Monotropism characterises autism by a default preference for monotropic processing. However, Priority Monotropism makes clearer claims about the nature of attention and attentional differences in autism by using the Priority Structure account of attention as a foundation. This allows it to have improved explanatory power over the extant account.

I have demonstrated this explanatory power by showing that Priority Monotropism can overcome both prongs of the heterogeneity problem. It overcomes the first prong of the heterogeneity problem by providing a conceptual link between the wide range of seemingly disparate symptoms and behaviours that are common to autism. I have shown that these symptoms and behaviours can all be explained by a default preference for monotropic processing (as it is understood by this revised account). The addition of the foundation from the Priority Structure account has allowed Priority Monotropism to provide an improved explanation of how the guidance of attention is influenced by monotropic processing. This allows for an improved understanding of why attention may be sticky and why autistic individuals may have an insistence for sameness and exhibit RRBIs. Priority Monotropism also overcomes the second prong of the heterogeneity problem by accounting for individual variance across symptom profiles. Priority Monotropism claims that the essence of autism is a specific way of attending, of having a narrow (spiky) priority structure; however, it does not claim that any specific contents necessarily need be contained within said priority structures. On this view, a preference for structuring priority structures into a specific shape profile (narrow/spiky) is what is common across autistic individuals. However, the specific priorities that autistic individuals hold are likely to vary and, as such, variance in symptom profiles is to be expected on this account.

In the final chapter of this thesis, I have also considered potential directions for future research regarding this theory. I have shown that Priority Monotropism may be able to explain why a wide range of seemingly unrelated conditions often co-occur with autism. I have also shown that Priority Monotropism has the potential to resolve inconsistencies in autism research findings and provide unifying explanations for these inconsistencies. Finally,

I have claimed that since Priority Monotropism has made clearer claims about the nature of attentional differences in autism and monotropic processing, it has laid a foundation for future examination into whether this account could be empirically testable. I have also claimed that providing clearer claims about monotropic processing may help to improve future support efforts. However, I have concluded that the examination of empirical testability and specific support measures is beyond the scope of this thesis.

References

- 'Access Consciousness' (2013) in Pashler, H. (ed.) *Encyclopedia of the Mind*. SAGE Publications. Available at: <https://doi.org/10.4135/9781452257044.n1> (Accessed: 21/04/2026).
- Adams, K., Smith, J., Brown, M., Bird, G., and Waite, P. (2025) 'Experiences of interoception and anxiety in autistic adolescents: A reflexive thematic analysis', *Autism*, 29(6), pp.1585-1596. Available at: <https://journals.sagepub.com/doi/10.1177/13623613251314595> (Accessed: 23/04/2026).
- Adams, K., Smith, J., Brown, M., Bird, G., and Waite, P. (2025) 'Experiences of interoception and anxiety in autistic adolescents: A reflexive thematic analysis'. *Autism*, 29(6), pp.1585–1596. Available at: <https://doi.org/10.1177/13623613251314595> (Accessed: 14/06/2026).
- Al Backer, N. (2015) 'Developmental regression in autism spectrum disorder', *Sudan J Paediatr.*, 15(1), pp. 21-26. Available at: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4949854/> (Accessed: 17/06/2024).
- Alkhaldi, R., Sheppard, E., Burdett, E., and Mitchell, P. (2021) 'Do Neurotypical People Like or Dislike Autistic People?', *Autism Adulthood.*, 3 (3), pp. 275-279. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8992906/> (Accessed: 02/10/2025).
- Allport, A., Antonis, B., and Reynolds, P. (1972) 'On the Division of Attention: A Disproof of the Single Channel Hypothesis', *Quarterly Journal of Experimental Psychology*, 24, pp. 225-235. Available at: <https://doi.org/10.1080/00335557243000102> (Accessed: 20/06/2026).
- American Psychiatric Association. (2000) *Diagnostic and statistical manual of mental disorders*. 4th edition. Washington, DC.
- American Psychiatric Association. (2013) *Diagnostic and statistical manual of mental disorders*. 5th edition. Washington, DC.
- Anderson, B. (2011) 'There is no Such Thing as Attention', *Front. Psychology*, 2. Available at: <https://doi.org/10.3389/fpsyg.2011.00246> (Accessed: 07/05/2024).
- Anderson, J. and Cushing, S. (2026) 'Introduction', in J. Anderson and S. Cushing (eds.), *Contemporary Philosophy of Autism*. Abingdon: Routledge. Available at: <https://doi-org.ezproxy1.lib.gla.ac.uk/10.4324/9781003627104> (Accessed: 28/03/2026).
- Annan, H. (2021) *Examining increased rates of synaesthesia in autism populations: the monotropic-compensatory account* [unpublished Master's dissertation]. University of Glasgow.
- Autism Alliance UK (n.d.) *Why we need to understand*. Available at: <https://www.autism-alliance.org.uk/news/why-we-need-to-understand> (Accessed: 14/09/2025).

- Autistic Self Advocacy Network (n.d.) *Identity-First Language*. Available at: <https://autisticadvocacy.org/about-asan/identity-first-language/> (Accessed: 06/06/2026).
- Autistica (2022) *Awareness isn't enough: understanding is critical to acceptance*. Available at: <https://www.autistica.org.uk/news/attitudes-index-news> (Accessed: 20/05/2024).
- Aylward, B., Gal-Szabo, D., and Taraman, S. (2021) 'Racial, Ethnic, and Sociodemographic Disparities in Diagnosis of Children with Autism Spectrum Disorder', *J Dev Behav Pediatr.*, 42(8), pp. 682-689. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8500365/> (Accessed: 06/06/2026).
- Banissy, M., Tester, V., Muggleton, N., Janik, A., Davenport, A., Franklin, A., Walsh, V. and Ward, J. (2013) 'Synesthesia for Color Is Linked to Improved Color Perception but Reduced Motion Perception', *Psychol Sci.*, vol.24, no.12 [online] available at <https://doi.org/10.1177/0956797613492424>
- Barahona-Corrêa, B. (2017) 'The High-Functioning Group: High-Functioning Autism and Asperger Syndrome in Adults', in B. Barahona-Corrêa and R. van der Gaag (eds.), *Autism Spectrum Disorders in Adults*. Springer. Available at: https://link.springer.com/chapter/10.1007/978-3-319-42713-3_5 (Accessed: 06/06/2026).
- Barned, N., Knapp, N., and Neuharth-Pritchett, S. (2011) 'Knowledge and Attitudes of Early Childhood Preservice Teachers Regarding the Inclusion of Children With Autism Spectrum Disorder', *Journal of Early Childhood Teacher Education*, 32(4), pp. 302-321. Available at: <https://doi.org/10.1080/10901027.2011.622235> (Accessed: 14/03/2026).
- Baron-Cohen, S. (1989) 'Perceptual role taking and protodeclarative pointing in autism', *British Journal of Developmental Psychology*, 7, pp. 113-127. Available at: <https://doi.org/10.1111/j.2044-835X.1989.tb00793.x> (Accessed: 15/05/2026).
- Baron-Cohen, S. (2000) 'Theory of mind and autism: A review', *Autism*, 23, pp. 169–184. Available at: [https://doi.org/10.1016/s0074-7750\(00\)80010-5](https://doi.org/10.1016/s0074-7750(00)80010-5) (Accessed: 05/06/2026).
- Baron-Cohen, S. (2002) 'The extreme male brain theory of autism', *Trends CognSci.*, 6 (6), pp. 248-254. Available at: <https://pubmed.ncbi.nlm.nih.gov/12039606/> (Accessed: 20/05/2026).
- Baron-Cohen, S. (2003) *The essential difference: Men, women and the extreme male brain*. London: Penguin.
- Baron-Cohen, S., Johnson, D., Asher, J., Wheelwright, S., Fisher, S., Gregersen, P., and Allison, C. (2013) 'Is synaesthesia more common in autism?', *Molecular Autism*, 4(40). Available at <https://doi.org/10.1186/2040-2392-4-40> (Accessed: 19/05/2026).
- Baron-Cohen, S., Leslie, A., and Frith, U. (1985) 'Does the autistic child have a "theory of mind" ?', *Cognition*, 21 (1), pp.37-46. Available at: <https://www.sciencedirect.com/science/article/pii/0010027785900228> (Accessed: 06/06/2026).

- Baron-Cohen, S., Robson, E., Lai, M., and Allison, C. (2016) 'Mirror-Touch Synaesthesia Is Not Associated with Heightened Empathy, and Can Occur with Autism', *PloS one*, 11(8). Available at: <https://doi.org/10.1371/journal.pone.0160543> (Accessed: 14/06/2026).
- Beccaria, F., Gagliardi, G., and Kissine, M. (2024) 'Atypical Regional Accent in Autistic Children: A Perception Study', *Autism Res.*, 18(2), pp.415-426. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11826007/> (Accessed: 03/12/2025).
- Behrmann, M. (2026) 'Visual agnosia in the era of behavioral and neural investigations', *Neuropsychologia*, 221. Available at: <https://doi.org/10.1016/j.neuropsychologia.2025.109299> (Accessed: 28/06/2026).
- Bellon-Harn, M., Wolford, G., Saar, K., Santhanam, S., and Hunt, M. (2026) 'Navigating "imaginary conversations" with neurotypicals: Insights from Instagram', *Research in Neurodiversity*, 2. Available at: <https://doi.org/10.1016/j.rin.2026.100020> (Accessed: 17/06/2026).
- Berlyne, D. (1974) 'Chapter 8 – ATTENTION', in E. Carterette and M. Friedman (eds.) *Historical and Philosophical Roots of Perception*, Academic Press, pp.123-147. Available at: <https://doi.org/10.1016/B978-0-12-161901-5.50015-7> (Accessed: 06/07/2026).
- Besag F. (2017) 'Epilepsy in patients with autism: links, risks and treatment challenges', *Neuropsychiatric disease and treatment*, 14. Available at: <https://doi.org/10.2147/NDT.S120509> (Accessed: 27/06/2026).
- Bettelheim, B. (1967) *The Empty Fortress: Infantile Autism and the Birth of the Self*. New York: The Free Press.
- Bogdashina, O. (2016) *Sensory Perceptual Issues in Autism and Asperger Syndrome*. 2nd ed. London: Jessica Kingsley Publishers.
- Booth, R. and Happé, F. (2010) "'Hunting with a knife and ... fork": Examining central coherence in autism, attention deficit/hyperactivity disorder, and typical development with a linguistic task', *Journal of Experimental Child Psychology*, 107 (4), pp. 337-395. Available at: <https://doi.org/10.1016/j.jecp.2010.06.003> (Accessed: 15/05/2026).
- Bradley, L., Shaw, R., Baron-Cohen, S., and Cassidy, S. (2021) 'Autistic Adults' Experiences of Camouflaging and Its Perceived Impact on Mental Health', *Autism in Adulthood.*, 3(4), pp. 320-329. Available at: <https://doi.org/10.1089/aut.2020.0071> (Accessed 06/06/2026).
- Braman, B., Brady, M., Linehan, S., and Williams, R. (1995) 'Facilitated Communication for Children with Autism: An Examination of Face Validity', *Behavioral Disorders*, 21(1), pp. 110-118. Available at: <https://www.jstor.org/stable/23888334> (Accessed: 06/06/2026).
- Bredemeier, K. and Simons, D. (2012) 'Working memory and inattention blindness', *Psychonomic Bulletin & Review*, 19, pp.239-244. Available at: <https://link.springer.com/article/10.3758/s13423-011-0204-8> (Accessed: 25/04/2025).

- British Medical Association (2024) *Autism spectrum disorder*. Available at: <https://www.bma.org.uk/what-we-do/population-health/improving-the-health-of-specific-groups/autism-spectrum-disorder> (Accessed: 10/06/2025).
- Bross, L. and Travers, J. (2017) 'Special Interest Areas and Employment Skills Programming for Secondary Students With Autism', *TEACHING Exceptional Children*, 50(2), pp.74-83. Available at: <https://doi.org/10.1177/0040059917730846> (Accessed: 15/06/2026).
- Carlsson, E., Miniscalco, C., Gillberg, C., and Johnels, J. (2018) 'Assessing False-Belief Understanding in Children with Autism Using a Computer Application: A Pilot Study', *Journal of psycholinguistic research*, 47 (5), pp. 1085-1099. Available at: <https://pubmed.ncbi.nlm.nih.gov/29582220/> (Accessed: 13/04/2026).
- Carpenter L. (2013) *DSM-5 autism spectrum disorder: guidelines criteria and exemplars*. Available at: <https://depts.washington.edu/dbpeds/Screening%20Tools/DSM-5%28ASD.Guidelines%29Feb2013.pdf> (Accessed: 01/07/2026).
- Carpita, B., Nardi, B., Tognini, V., Poli, F., Amatori, G., Cremone, I., Pini, S., and Dell'Osso, L. (2024) 'Autistic Traits and Somatic Symptom Disorders: What Is the Link?', *Brain sciences*, 14(3). Available at: <https://doi.org/10.3390/brainsci14030274> (Accessed: 17/06/2026).
- Casey, B., Gordon, C., Mannheim, G., and Rumsey, J. (1993) 'Dysfunctional attention in autistic savants', *Journal of clinical and experimental neuropsychology*, 15(6), pp.933-946. Available at: <https://doi.org/10.1080/01688639308402609> (Accessed: 10/06/2026).
- Cassidy, S., Hannant, P., Tavassoli, T., Allison, C., Smith, P., and Baron-Cohen, S. (2016) 'Dyspraxia and autistic traits in adults with and without autism spectrum conditions', *Molecular autism*, 7(48). Available at: <https://doi.org/10.1186/s13229-016-0112-x> (Accessed: 12/06/2026).
- Center for Philosophy of Science (2019) *Do We Know What Attention Is: Wayne Wu ALS Talk, Dec. 6 2019*. [online video] Available at: <https://www.youtube.com/watch?v=MCaYTcpoMnQ> (Accessed: 12/07/2024).
- Centers for Disease Control and Prevention (2024) *Signs and Symptoms of Autism Spectrum Disorder*. Available at: <https://www.cdc.gov/autism/signs-symptoms/index.html> (Accessed: 10/06/2026).
- Chapman, R. (2020) 'The reality of autism: On the metaphysics of disorder and diversity', *Philosophical Psychology*, 33(6), pp. 799-819. Available at: <https://doi.org/10.1080/09515089.2020.1751103> (Accessed: 04/09/2025).
- Children's Hospital of Philadelphia Research Institute (2020) *Diagnostic Criteria for Autism Spectrum Disorder in the DSM-5*. Available at: <https://www.research.chop.edu/car-autism-roadmap/diagnostic-criteria-for-autism-spectrum-disorder-in-the-dsm-5> (Accessed: 07/04/2026).

Chung, U. and Kim, J. (2024) 'Common Comorbid Condition of Patients With Autism Spectrum Disorder and Pharmacotherapy for Patients With Autism Spectrum Disorder', *Journal of the Korean Academy of Child and Adolescent Psychiatry*, 35(1), pp.39–43. Available at: <https://doi.org/10.5765/jkacap.230006> (Accessed: 20/06/2026).

Conner, C., Ionadi, A., Mazefsky, C. (2023) 'Recent Research Points to a Clear Conclusion: Autistic People are Thinking About, and Dying by, Suicide at High Rates', *Pa J Posit Approaches*, 12(3), pp. 69-76. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11042491/> (Accessed: 06/06/2026).

Corbett, J., Venuti, P., and Melcher, D. (2016) 'Perceptual Averaging in Individuals with Autism Spectrum Disorder', *Front. Psychol.*, 7. Available at: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2016.01735/full> (Accessed: 21/04/2026).

Cribb, S., Olaithe, M., Di Lorenzo, R., Dunlop, P., and Maybery, M. (2016) 'Embedded Figures Test Performance in the Broader Autism Phenotype: A Meta-analysis', *J Autism Dev Disord.*, 46(9), pp.2924-39. Available at: <https://pubmed.ncbi.nlm.nih.gov/27312717/> (Accessed: 14/03/2023).

Cristofori, I., Cohen-Zimmerman, S., and Grafman, J. (2019) 'Executive functions', *Handbook of clinical neurology*, 163, pp. 197-219. Available at: <https://doi.org/10.1016/B978-0-12-804281-6.00011-2> (Accessed: 02/06/2026).

Cushing, S. (2012) 'Autism: The Very Idea', in J. Anderson and S. Cushing (eds.), *The Philosophy of Autism*. Plymouth: Rowman & Littlefield Publishers, Inc. Available at: <https://ebookcentral.proquest.com/lib/gla/detail.action?docID=1318925> (Accessed: 13/12/2025).

Cushing, S. (2026) 'The Thing of it Isn't', in J. Anderson and S. Cushing (eds.), *Contemporary Philosophy of Autism*. Abingdon: Routledge. Available at: <https://doi-org.ezproxy1.lib.gla.ac.uk/10.4324/9781003627104> (Accessed: 28/03/2026).

de Jager, P. and Condry, J. (2020) 'Weak central coherence is a syndrome of autism spectrum disorder during teacher–learner task instructions', *South African Journal of Childhood Education*, 10 (1). Available at: <https://doi.org/10.4102/sajce.v10i1.785> (Accessed: 10/06/2026).

Department for Work & Pensions (2024) *The Buckland Review of Autism Employment: report and recommendations*. Available at: <https://www.gov.uk/government/publications/the-buckland-review-of-autism-employment-report-and-recommendations/the-buckland-review-of-autism-employment-report-and-recommendations> (Accessed: 12/07/2025).

Digital Science (2026) *Dimensions* [Online]. Available at: www.dimensions.ai (Accessed: 10/06/2026).

DiGregorio, N. and Sharma, S. (2023) 'Physiology, Secretin', In: *StatPearls* [Internet]. StatPearls Publishing. Available at: <https://www.ncbi.nlm.nih.gov/books/NBK537116/> (Accessed: 18/02/2026).

D'Souza, R. and Hooten, W. (2023) 'Somatic Symptom Disorder', in *StatPearls*. StatPearls Publishing. Available at: <https://www.ncbi.nlm.nih.gov/books/NBK532253/>

Durgin, F., Tripathy, S., and Levi, D. (1995) 'On the Filling in of the Visual Blind Spot: Some Rules of Thumb', *Perception.*, 24(7). Available at: <https://doi.org/10.1068/p240827> (Accessed: 20/04/2025).

Dyspraxia UK (n.d.) *Diagnostic Criteria*. Available at: <https://dyspraxiauk.com/diagnostic-criteria> (Accessed: 12/06/2026).

Erez, A., Kizony, R., Shahar, M., and Katz, N. (2006) 'Visual spatial search task (VISSTA): a computerized assessment and training program', *Proc. 6th Intl Conf. Disability, Virtual Reality & Assoc. Tech., Esbjerg, Denmark, 2006.*, Hebrew University & Hadassah, Jerusalem/University of Haifa. Available at https://www.academia.edu/12421731/Proceedings_of_the_6th_International_Conference_on_Computational_Methods_in_Systems_Biology (Accessed: 22/01/2026).

Falter-Wagner, C., Bloch, C., Robles, M., Horch, L., Vogeley, K., and Georgescu, A. (2022) 'Figure-Disembedding Is Inferior in Non-autistic Compared to Autistic Individuals but Can Be Improved by Training', *Front. Psychol.*, 13. Available at: <https://doi.org/10.3389/fpsyg.2022.857630> (Accessed: 14/03/2023).

Fletcher-Watson, S. and Happé, F. (2019) *Autism: A New Introduction to Psychological Theory and Current Debate*. Routledge. Available at: <https://doi-org.ezproxy1.lib.gla.ac.uk/10.4324/9781315101699> (Accessed: 04/09/2025).

Furlan, M. and Smith, A. (2016) 'Global Motion Processing in Human Visual Cortical Areas V2 and V3', *J Neurosci.*, 36(27), pp.7314-24. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6705536/> (Accessed: 20/03/2026).

Gabis, L., Attia, O., Goldman, M., Barak, N., Tefera, P., Shefer, S., Shaham, M., and Lerman-Sagie, T. (2022) 'The myth of vaccination and autism spectrum', *Eur J Paediatr Neurol.*, 36, pp. 151-158. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8694782/> (Accessed: 19/07/2025).

García, F. (2026) 'Autism from the Second Person Perspective', in J. Anderson and S. Cushing (eds.), *Contemporary Philosophy of Autism*. Abingdon: Routledge. Available at: <https://doi-org.ezproxy1.lib.gla.ac.uk/10.4324/9781003627104> (Accessed: 27/03/2026).

Gernsbacher, M. and Yergeau, M. (2019) 'Empirical failures of the claim that autistic people lack a theory of mind', *Archives of Scientific Psychology*, 7(1), pp.102–118. Available at: <https://doi.org/10.1037/arc0000067> (Accessed: 27/06/2023).

Glenwright, M., Scott, R., Bilevicius, E., Pronovost, M., and Hanlon-Dearman, A. (2021) 'Children With Autism Spectrum Disorder Can Attribute False Beliefs in a Spontaneous-

Response Preferential-Looking Task', *Front. Commun.*, 6. Available at: <https://doi.org/10.3389/fcomm.2021.669985> (Accessed: 14/05/2026).

Gomez, R. and Sheikh, S. (2023) *Opening opportunities: Improving employment prospects for autistic people*. Available at: <https://www.autistica.org.uk/downloads/files/PBE-improve-autistic-employment-report-Autistica.pdf> (Accessed: 06/06/2026).

Grandin, T. (2008) *The Way I See It: A Personal Look at Autism and Asperger's*. Arlington: Future Horizons.

Grandin, T. and Scariano, M. (1986) *Emergence, Labelled Autistic*. Novato, CA: Arena.

Grelotti, D., Klin, A., Gauthier, I., Skudlarski, P., Cohen, D., Gore, J., Volkmar, F., and Schultz, R. (2005) 'fMRI activation of the fusiform gyrus and amygdala to cartoon characters but not to faces in a boy with autism', *Neuropsychologia*, 43(3), pp.373-385. Available at: <https://www.sciencedirect.com/science/article/abs/pii/S0028393204001575> (Accessed: 12/02/2024).

Guan, J., Blanchard, A., DiGuseppi, C., Chihuri, S., and Li, G. (2021) 'Homicide Incidents Involving Children with Autism Spectrum Disorder as Victims Reported in the US News Media, 2000–2019', *J Autism Dev Disord.*, 52(4), pp. 1673-1677. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10124270/> (Accessed: 06/06/2026).

Guy, J., Mottron, L., Berthiaume, C., and Bertone, A. (2019) 'A Developmental Perspective of Global and Local Visual Perception in Autism Spectrum Disorder', *Journal of Autism and Developmental Disorders*, 49. Available at <https://doi.org/10.1007/s10803-016-2834-1> (Accessed: 22/01/2026).

Happé F. (2013 a) 'Embedded Figures Test (EFT)'. in: Volkmar F.R. (eds.) *Encyclopedia of Autism Spectrum Disorders*. Springer. Available at: https://doi.org/10.1007/978-1-4419-1698-3_1726 (Accessed: 22/01/2026).

Happé, F. (2013 b) 'Weak Central Coherence', in F. Volkmar (ed), *Encyclopedia of Autism Spectrum Disorders*. Springer, New York. Available at: https://doi.org/10.1007/978-1-4419-1698-3_1744 (Accessed: 20/04/2026).

Happé, F. and Frith, U. (2006) 'The weak coherence account: detail-focused cognitive style in autism spectrum disorders', *J Autism Dev Disord.*, 36 (1), pp. 5-25. Available at: <https://pubmed.ncbi.nlm.nih.gov/16450045/> (Accessed: 04/03/2026).

Hill, E. (2004) 'Evaluating the theory of executive dysfunction in autism', *Developmental Review*, 24 (2), pp. 189-233. Available at: <https://doi.org/10.1016/j.dr.2004.01.001> (Accessed: 02/06/2026).

Hogeveen, J. and Grafman, J. (2021) 'Alexithymia', *Handbook of clinical neurology*, 183, pp.47–62. Available at: <https://doi.org/10.1016/B978-0-12-822290-4.00004-9>

- Hommel, B., Chapman, C., Cisek, P., Neyedli, H., Song, J., and Welsh, T. (2019) 'No one knows what attention is', *Atten Percept Psychophys.*, 81, pp. 2288-2303. Available at: <https://doi.org/10.3758/s13414-019-01846-w> (Accessed: 14/06/2024).
- Horvath, K., Stefanatos, G., Sokolski, K., Wachtel, R., Nabors, L., and Tildon, J. (1998) 'Improved social and language skills after secretin administration in patients with autistic spectrum disorders', *J Assoc Acad Minor Phys.*, 9(1), pp. 9-15. Available at: <https://pubmed.ncbi.nlm.nih.gov/9585670/> (Accessed: 12/01/2026).
- Howlin, P., Goode, S., Hutton, J., and Rutter, M. (2009) 'Savant skills in autism: psychometric approaches and parental reports', *Philosophical transactions of the Royal Society of London, Series B, Biological sciences*, 364(1522), pp.1359–1367. Available at: <https://doi.org/10.1098/rstb.2008.0328> (Accessed: 11/06/2026).
- Huang, W., Leventhal, B., Lin, C., and Chien, Y. (2025) 'Health anxiety and somatic symptoms in adults on the autism spectrum', *Psychiatry Research*, 350. Available at: <https://doi.org/10.1016/j.psychres.2025.116550> (Accessed: 12/06/2026).
- Hughes, J., Ward, J., Gruffydd, E., Baron-Cohen, S., Smith, P., Allison, C., and Simner, J. (2018) 'Savant syndrome has a distinct psychological profile in autism', *Molecular autism*, 9(53). Available at: <https://doi.org/10.1186/s13229-018-0237-1> (Accessed: 12/06/2026).
- Hus, Y. and Segal, O. (2021) 'Challenges Surrounding the Diagnosis of Autism in Children', *Neuropsychiatr Dis Treat.*, 3(17), pp. 3509-3529. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8654688/> (Accessed: 06/06/2026).
- Isaacs, P. (2026) *Visual perception in autism*. Available at: <https://www.autism.org.uk/learn/knowledge-hub/professional-practice/visual-perception> (Accessed: 01/07/2026).
- James, S., Stevenson, S., Silove, N., and Williams, K. (2015) 'Chelation for autism spectrum disorder (ASD)', *Cochrane Database Syst Rev.*, 5(5). Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6457964/> (Accessed: 12/01/2026).
- James, W. (1890) *The Principles of Psychology*. New York: Dover Publications.
- Jensen, M., Yao, R., Street, W., and Simons, D. (2011) 'Change blindness and inattentional blindness', *WIREs Cogn Sci*, 2, pp. 529-546. Available at: <https://doi.org/10.1002/wcs.130> (Accessed: 21/04/2025).
- Ji, L., Yap, S., Best, M., and McGeorge, K. (2019) 'Global Processing Makes People Happier Than Local Processing', *Front. Psychol.*, 10. Available at: <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2019.00670/full> (Accessed: 12/04/2026).
- Johnson, A. and Proctor, R. (2003) *Attention: Theory and Practice*. SAGE Publications. Available at: <https://ebookcentral.proquest.com/lib/gla/detail.action?docID=1920463> (Accessed: 25/01/2025).

- Kaldy, Z., Giserman, I., Carter, A. and Blaser, E. (2016) 'The Mechanisms Underlying the ASD Advantage in Visual Search', *J Autism Dev Disord.*, 46(5). Available at <https://doi.org/10.1007/s10803-013-1957-x>
- Kalingel-Levi, M., Schreuer, N., Granovsky, Y., Bar-Shalita, T., Weissman-Fogel, I., Hoffman, T., and Gal, E. (2022) "“When I’m in Pain, Everything Is Overwhelming”": Implications of Pain in Adults With Autism on Their Daily Living and Participation', *Front. Psychol.*, 13. Available at: <https://doi.org/10.3389/fpsyg.2022.911756> (Accessed: 12/03/2024).
- Kantawala, B., Abu-Bakr, A., Kasini, B., Ndayambaje, M., Ian Soh, S., Nazir, A., Wojtara, M., and Uwishema, O. (2023) 'Exploring the landscape of autism in Africa: challenges in diagnosis, support, and resources – a short communication', *Ann Med Surg (Lond)*, 85(11), pp. 5838-5841. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10617929/> (Accessed: 06/06/2026).
- Katusic, M., Myers, S., Weaver, A., and Voigt, R. (2021) 'IQ in Autism Spectrum Disorder: A Population-Based Birth Cohort Study', *Pediatrics*, 148(6). Available at: <https://doi.org/10.1542/peds.2020-049899> (Accessed: 10/04/2025).
- Kawczynska, C. (2009) 'Q&A with Temple Grandin', *The Bark*. Available at: <https://thebark.com/content/qatemple-grandin> (Accessed: 25/04/2023).
- Kenny, L., Hattersley, C., Molins, B., Buckley, C., Povey, C., and Pellicano E. (2015) 'Which terms should be used to describe autism? Perspectives from the UK autism community', *Autism*, 20(4). Available at: <https://journals.sagepub.com/doi/10.1177/1362361315588200> (Accessed: 04/06/2026).
- Khatab, S., Hijab, M., Othman, A., and Al-Thani, D. (2024) 'Collaborative play for autistic children: A systematic literature review', *Entertainment Computing*, 50. Available at: <https://doi.org/10.1016/j.entcom.2024.100653> (Accessed: 15/04/2025).
- Kinnaird, E., Norton, C., Pimblett, C., Stewart, C., and Tchanturia, K. (2019) 'Eating as an autistic adult: An exploratory qualitative study', *PLoS One*, 14(8). Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6715205/> (Accessed: 12/04/2025).
- Kinnaird, E., Stewart, C., and Tchanturia, K. (2019) 'Investigating alexithymia in autism: A systematic review and meta-analysis', *European psychiatry*, 55, pp.80–89. Available at: <https://doi.org/10.1016/j.eurpsy.2018.09.004> (Accessed: 11/06/2026).
- Klin, A., Jones, W., Schultz, R., and Volkmar, F. (2003) 'The enactive mind, or from actions to cognition: lessons from autism', *Philos Trans R Soc Lond B Biol Sci.*, 358(1430), pp. 345-360. Available at: <https://doi.org/10.1098/rstb.2002.1202> (Accessed: 12/07/2023).
- Koldewyn, K., Jiang, Y., Weigelt, S., and Kanwisher, N. (2013) 'Global/local processing in autism: not a disability, but a disinclination', *J Autism Dev Disord.*, 43 (10). Available at: <https://pubmed.ncbi.nlm.nih.gov/23378063/> (Accessed: 28/04/2026).

- Krahn, T. and Fenton, A. (2012) 'The Extreme Male Brain Theory of Autism and the Potential Adverse Effects for Boys and Girls with Autism', *Journal of Bioethical Inquiry*, 9, pp. 93-103. Available at: <https://doi.org/10.1007/s11673-011-9350-y> (Accessed: 10/06/2025).
- Lane, R., Solms, M., Weihs, K., Hishaw, A., and Smith, R. (2020) 'Affective agnosia: a core affective processing deficit in the alexithymia spectrum', *BioPsychoSocial Med*, 14(20). Available at: <https://doi.org/10.1186/s13030-020-00184-w> (Accessed: 23/06/2026).
- Lavie, N. (1995) 'Perceptual load as a necessary condition for selective attention', *Journal of Experimental Psychology: Human Perception and Performance*, 21(3), pp.451–468. Available at: <https://doi.org/10.1037/0096-1523.21.3.451> (Accessed: 06/07/2026).
- Leask, J. and Silove, N. (2005) 'Evidence for autism in folklore?', *Arch Dis Child*, 90(3). Available at: <https://pubmed.ncbi.nlm.nih.gov/15723914/> (Accessed: 18/05/2026).
- Lindsay, G. (2020) 'Attention in Psychology, Neuroscience, and Machine Learning', *Front Comput Neurosci*, 14(29). Available at: <https://www.frontiersin.org/journals/computational-neuroscience/articles/10.3389/fncom.2020.00029/full> (Accessed: 14/06/2024).
- Long, R. (2025) 'Access Points: Understanding Special Interests Through Autistic Narratives', *Autism Adulthood*, 7(1), pp.100-111. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11937755/> (Accessed: 20/03/2026).
- Loomes, R., Hull, L., and Mandy, W. (2017) 'What Is the Male-to-Female Ratio in Autism Spectrum Disorder? A Systematic Review and Meta-Analysis', *Journal of the American Academy of Child & Adolescent Psychiatry*, 56(6), pp. 466-474. Available at: <https://doi.org/10.1016/j.jaac.2017.03.013> (Accessed: 06/06/2026).
- Lordan, R., Storni, C., and De Benedictis, C. (2021) 'Autism Spectrum Disorders: Diagnosis and Treatment', in A. Grubucker (ed.), *Autism Spectrum Disorders*. Exon Publications. Available at: <https://www.ncbi.nlm.nih.gov/books/NBK573609/> (Accessed: 15/07/2025).
- Loughran, E. (2020) *Why autism research needs more input from autistic people*. Available at: <https://www.spectrumnews.org/opinion/viewpoint/why-autismresearch-needs-more-input-from-autistic-people/> (Accessed: 10/03/2023).
- MacLennan, K., O'Brien, S., and Tavassoli, T. (2022) 'In Our Own Words: The Complex Sensory Experiences of Autistic Adults', *J Autism Dev Disord*, 52, pp. 3061–3075. Available at: <https://doi.org/10.1007/s10803-021-05186-3> (Accessed: 24/07/2023).
- Marco, E., Hinkley, L., Hill, S., and Nagarajan, S. (2011) 'Sensory Processing in Autism: A Review of Neurophysiologic Findings', *Pediatric Research*, 69, pp. 48-54. Available at: <https://doi.org/10.1203/pdr.0b013e3182130c54> (Accessed: 10/03/2023).
- Mole, C. (2010) *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford: Oxford University Press.
- Morris, I., Sykes, J., Paulus, E., Dameh, A., Razzaque, A., Esch, L., Gruenig, J., and Zelazo, P. (2025) 'Beyond self-regulation: Autistic experiences and perceptions of stimming',

Neurodiversity, 3. Available at: <https://doi.org/10.1177/27546330241311096> (Accessed: 15/02/2026).

Mottron, L., Dawson, M., Soulières, I., Hubert, B., and Burack, J. (2006) 'Enhanced Perceptual Functioning in Autism: An Update, and Eight Principles of Autistic Perception', *Journal of Autism and Developmental Disorders*, 36. Available at: <https://doi.org/10.1007/s10803-005-0040-7> (Accessed: 20/04/2026).

Mukhopadhyay, T. (2008) *How Can I Talk If My Lips Don't Move? Inside My Autistic Mind*. New York: Arcade Publishing.

Murray, D. (2018) 'Monotropism – An Interest Based Account of Autism', *Encyclopedia of Autism Spectrum Disorders*. Available at: https://doi.org/10.1007/978-1-4614-6435-8_102269-1 (Accessed: 10/03/2023).

Murray, D., Lesser, M., and Lawson, W. (2005) 'Attention, monotropism and the diagnostic criteria for autism', *Autism*, 9(2), pp. 139–156. Available at: <https://doi.org/10.1177/1362361305051398> (Accessed: 19/02/2023).

Nah, Y. and Poon, K. (2011) 'The perception of social situations by children with autism spectrum disorders', *Autism*, 15(2), pp.185-203. Available at: <https://journals.sagepub.com/doi/10.1177/1362361309353616> (Accessed: 25/03/2024).

National Autistic Society (2021) *National Autistic Society Education Report*. Available at: <https://www.nen.org.uk/resources/national-autistic-society-education-report> (Accessed: 20/05/2026).

National Autistic Society (n.d. a) *What is autism?* Available at: <https://www.autism.org.uk/advice-and-guidance/what-is-autism> (Accessed: 14/03/2026).

National Autistic Society (n.d. b) *Criteria and tools use in an autism assessment*. Available at: <https://www.autism.org.uk/advice-and-guidance/diagnosis/assessment-and-diagnosis/criteria-and-tools-used-in-an-autism-assessment> (Accessed: 10/08/2025).

National Autistic Society (n.d. c) *Autistic women and girls*. Available at: <https://www.autism.org.uk/advice-and-guidance/identity/autistic-women-and-girls> (Accessed: 17/03/2026).

National Autistic Society (n.d. d) *Masking*. Available at: <https://www.autism.org.uk/advice-and-guidance/behaviour/masking> (Accessed: 12/07/2025).

National Autistic Society (n.d. e) *Signs that a child or adult may be autistic*. Available at: <https://www.autism.org.uk/advice-and-guidance/diagnosis/before-diagnosis/signs-that-a-child-or-adult-may-be-autistic> (Accessed: 10/06/2026).

National Autistic Society (n.d. f) *Autism and sensory processing*. Available at: <https://www.autism.org.uk/advice-and-guidance/about-autism/sensory-processing> (Accessed: 21/04/2026).

National Autistic Society (n.d. g) *Meltdowns – a guide for all audiences*. Available at: <https://www.autism.org.uk/advice-and-guidance/behaviour/meltdowns/all-audiences> (Accessed: 10/06/2026).

National Autistic Society (n.d. h) *Autistic Children and Motor Skills*. Available at: <https://www.autism.org.uk/learn/knowledge-hub/professional-practice/motor-skills> (Accessed: 10/06/2026).

National Autistic Society (n.d. i) *Repeated movements and behaviour (stimming)*. Available at: <https://www.autism.org.uk/advice-and-guidance/about-autism/repeated-movements-and-behaviour-stimming> (Accessed: 12/05/2026).

National Autistic Society (n.d. j) *Focused and dedicated interests*. Available at: <https://www.autism.org.uk/advice-and-guidance/about-autism/focused-and-dedicated-interests> (Accessed: 20/06/2025).

National Autistic Society (n.d. k) *Related conditions - a guide for all audiences*. Available at: <https://www.autism.org.uk/advice-and-guidance/related-conditions/related-conditions/all-audiences> (Accessed: 18/05/2026).

National Autistic Society (n.d. l) *Autism and catastrophising*. Available at: <https://www.autism.org.uk/learn/knowledge-hub/professional-practice/autism-catastrophising> (Accessed: 07/06/2026).

National Center for Biotechnology Information (2026) *Autism search results trend graph*, *PubMed*. Available at: [nih.gov](https://pubmed.ncbi.nlm.nih.gov/11257286/) (Accessed: 25/01/2026).

National Health Service (2025) *What is autism?* Available at: <https://www.nhs.uk/conditions/autism/what-is-autism/> (Accessed: 04/02/2026).

Neufeld, J., Sinke, C., Zedler, M., Emrich, H., and Szycik, G. (2012) ‘Reduced audio–visual integration in synaesthetes indicated by the double-flash illusion’, *Brain Research*, 1473. Available at: <https://doi.org/10.1016/j.brainres.2012.07.011> (Accessed: 12/06/2026).

Norman, L., Heywood, C., and Kentridge, R. (2013) ‘Object-based attention without awareness’, *Psychological science*, 24(6), pp. 836–843. Available at: <https://doi.org/10.1177/0956797612461449> (Accessed: 04/06/2026).

Offit, P. (2010) *Autism’s False Prophets*. New York: Columbia University Press.

Olk, B., Harvey, M., Dow, L., and Murphy, P. (2001) ‘Illusion processing in hemispatial neglect’, *Neuropsychologia*, 39(6). Available at: <https://pubmed.ncbi.nlm.nih.gov/11257286/> (Accessed: 20/06/2026).

Organization for Autism Research (2016) ‘*But She’s Too Social To Be Autistic*’ – *WRONG!*. Available at: <https://researchautism.org/blog/but-shes-too-social-to-be-autistic-wrong/> (Accessed: 24/08/2024).

Ostrolenk, A., Boisvert, M., and Motttron, L. (2025) ‘What do autistic children who are interested in letters and numbers do with them? A qualitative study’, *International Journal of*

Qualitative Studies on Health and Well-being, 20(1). Available at:

<https://doi.org/10.1080/17482631.2025.2500851> (Accessed: 20/06/2026).

Ostrolenk, A., Forgeot d'Arc, B., Jelenic, P., Samson, F., and Mottron, L. (2017)

'Hyperlexia: Systematic review, neurocognitive modelling, and outcome', *Neuroscience & Biobehavioral Reviews*, 79, pp.134-149. Available at:

<https://doi.org/10.1016/j.neubiorev.2017.04.029> (Accessed: 18/06/2026).

Ouhtit, A., Al-Farsi, Y., Al-Sharbati, M., Waly, M., Gupta, I., Al-Farsi, O., Al-Khaduri, M., Al-Shafae, M., and Al-Adawi, S. (2015) 'Underlying Factors Behind the Low Prevalence of Autism Spectrum Disorders in Oman', *Sultan Qaboos Univ Med J.*, 15(2), pp. 213-217.

Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4450784/> (Accessed: 06/06/2026).

Overweg, J., Hartman, C., and Hendriks, P. (2018) 'Children with autism spectrum disorder show pronoun reversals in interpretation', *Journal of abnormal psychology*, 127(2), pp.228–238. Available at: <https://doi.org/10.1037/abn0000338> (Accessed: 13/06/2026).

Peacocke, C. (1998) 'Conscious attitudes, attention and self knowledge', in C. McDonald, B. Smith, and C. Wright (eds.), *Knowing Our Own Minds: Essays on Self Knowledge*. Oxford: Oxford University Press.

Pellicano, E. and Burr, D. (2012) 'When the world becomes 'too real': a Bayesian explanation of autistic perception', *Trends in Cognitive Sciences*, 16, pp. 504-510. Available at: <https://doi.org/10.1016/j.tics.2012.08.009> (Accessed: 06/07/2026).

Pillsbury, W. (1908) *Attention*. [Reprinted in 1973]. New York: Arno Press.

Polyanskaya, I., Eigsti, I., Brauner, T., and Blackburn, P. (2021) 'Second-order false beliefs and linguistic recursion in autism spectrum disorder', *J Autism Dev Disord.*, 52 (9), pp. 3991-4006. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8920946/> (Accessed: 10/05/2026).

Preece, D., Becerra, R., Allan, A., Robinson, K., and Dandy, J. (2017) 'Establishing the theoretical components of alexithymia via factor analysis: Introduction and validation of the attention-appraisal model of alexithymia', *Personality and Individual Differences*, 119, pp.341-352. Available at: <https://doi.org/10.1016/j.paid.2017.08.003> (Accessed: 12/06/2026).

Prizant, B., Wetherby, A., Rubin, E., Laurent, A., and Rydell, P. (2006) *The SCRETS Model*. Vol 1 Assessment; Vol 2 Program Planning and Intervention. Baltimore, MD: Paul H. Brookes.

Raman, R. and Sarkar, S. (2016) 'Predictive Coding: A Possible Explanation of Filling-In at the Blind Spot', *PLoS One.*, 11(3). Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4784844/> (Accessed: 21/04/2025).

Ramos-Cabo, S., Vulchanov, V. and Vulchanova, M. (2021) 'Different Ways of Making a Point: A Study of Gestural Communication in Typical and Atypical Early Development', *Autism Research*, 14, pp. 984-996. Available at: <https://doi.org/10.1002/aur.2438> (Accessed: 18/05/2026).

- Rao, T. and Andrade, C. (2011) 'The MMR vaccine and autism: Sensation, refutation, retraction, and fraud', *Indian J Psychiatry*, 53(2), pp. 95-96. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC3136032/> (Accessed: 14/03/2026).
- Rigby, S., Stoesz, B. and Jakobson, L. (2016) 'Gaze patterns during scene processing in typical adults and adults with autism spectrum disorders', *Research in Autism Spectrum Disorders*, 25, pp. 24-36. Available at: <https://www.sciencedirect.com/science/article/pii/S1750946716300113> (Accessed: 01/05/2024).
- Ro, T. and Rafal, R. (1996) 'Perception of geometric illusions in hemispatial neglect', *Neuropsychologia*, 34(10), pp.973-978. Available at: <https://rolab.ws.gc.cuny.edu/files/2017/05/Ro96.pdf> (Accessed: 20/06/2026).
- Russell, G., Stapley, S., Newlove-Delgado, T., Salmon, A., White, R., Warren, F., Pearson, A., and Ford, T. (2021) 'Time trends in autism diagnosis over 20 years: a UK population-based cohort study', *The Journal of Child Psychology and Psychiatry*, 63, pp. 674-682. Available at: <https://doi.org/10.1111/jcpp.13505> (Accessed: 14/03/2026).
- Sample, R. (2012) 'Autism and the Extreme Male Brain', in J. Anderson and S. Cushing (eds.), *The Philosophy of Autism*. Plymouth: Rowman & Littlefield Publishers, Inc.
- Sanders, A. and Neumann, O. (1996) 'Introduction', in O. Neumann and A. Sanders (eds.), *Handbook of Perception and Action*. Academic Press. Available at: [https://doi.org/10.1016/S1874-5822\(96\)80017-X](https://doi.org/10.1016/S1874-5822(96)80017-X) (Accessed: 06/07/2026).
- Sandler, A., Sutton, K., DeWesse, J., Girardi, M., Sheppard, V., and Bodfish, J. (1999) 'Lack of Benefit of a Single Dose of Synthetic Human Secretin in the Treatment of Autism and Pervasive Developmental Disorder', *N Engl J Med*, 341(24), pp. 1801-1806. Available at: <https://www.nejm.org/doi/full/10.1056/NEJM199912093412404> (Accessed: 20/04/2026).
- Sanz-Cervera, P., Fernández-Andrés, M., Pastor-Cerezuela, G., and Tárraga-Mínguez, R. (2017) 'Pre-Service Teachers' Knowledge, Misconceptions and Gaps About Autism Spectrum Disorder', *Teacher Education and Special Education: The Journal of the Teacher Education Division of the Council for Exceptional Children*, 40(3), pp. 212-224. Available at: <https://journals.sagepub.com/doi/full/10.1177/0888406417700963> (Accessed: 06/06/2026).
- Schaaf, R. (2020) 'Sensory Features as a Marker of Autism Spectrum Disorders', *Encyclopedia of Autism Spectrum Disorders*. Available at: https://doi.org/10.1007/978-1-4614-6435-8_102455-1 (Accessed: 18/04/2023).
- Severson, K., Aune, J., and Jodlowski, D. (2008) 'Bruno Bettelheim, Autism, and the Rhetoric of Scientific Authority', in M. Osteen (ed.), *Autism and Representation*. Routledge. Available at: <https://www.taylorfrancis.com/chapters/edit/10.4324/9780203935088-8/bruno-bettelheim-autism-rhetoric-scientific-authority-katherine-demaria-severson-james-arnt-aune> (Accessed: 20/04/2026).

- Shalev, I., Warrier, V., Greenberg, D., Smith, P., Allison, C., Baron-Cohen, S., Eran, A., and Uzefovsky, F. (2022) 'Reexamining empathy in autism: Empathic disequilibrium as a novel predictor of autism diagnosis and autistic traits', *Autism research : official journal of the International Society for Autism Research*, 15 (10), pp. 1917-1928. Available at: <https://doi.org/10.1002/aur.2794> (Accessed: 07/06/2026).
- Simner, J., Mayo, N., and Spiller, M. (2009) 'A foundation for savantism? Visuo-spatial synaesthetes present with cognitive benefits', *Cortex*, 45(10), pp.1246-1260, Available at: <https://doi.org/10.1016/j.cortex.2009.07.007> (Accessed: 12/06/2026).
- Simner, J., Mulvenna, C., Sagiv, N., Tsakanikos, E., Witherby, S., Fraser, C., Scott, K. and Ward, J. (2006) 'Synaesthesia: The prevalence of atypical cross-modal experiences', *Perception*, 35(8). Available at <https://doi.org/10.1068/p5469> (Accessed: 20/05/2026).
- Simons, D. and Chabris, C. (1999) 'Gorillas in our midst: sustained inattention blindness for dynamic events', *Perception.*, 28(9), pp.1059-74. Available at: <https://pubmed.ncbi.nlm.nih.gov/10694957/> (Accessed: 25/04/2025).
- Sinke, C., Neufeld, J., Zedler, M., Emrich, H., Bleich, S., Münte, T., and Szycik, G. (2012) 'Reduced audiovisual integration in synesthesia – evidence from bimodal speech perception', *Journal of Neuropsychology*, 8(1). Available at <https://doi.org/10.1111/jnp.12006> (Accessed: 12/06/2026).
- Solazzo, S., Kojovic, N., Robain, F., and Schaer, M. (2021) 'Measuring the Emergence of Specific Abilities in Young Children with Autism Spectrum Disorders: The Example of Early Hyperlexic Traits', *Brain sciences*, 11(6). Available at: <https://doi.org/10.3390/brainsci11060692> (Accessed: 20/05/2026).
- Sommer, M., Döhl, K., Jarvers, I., Blaas, L., Singer, M., Nöth, V., Schuwerk, T., and Rupprecht, R. (2018) 'False Belief Reasoning in Adults with and without Autistic Spectrum Disorder: Similarities and Differences', *Frontiers in psychology*, 9. Available at: <https://doi.org/10.3389/fpsyg.2018.00183> (Accessed: 10/05/2026).
- Soydaner, D. (2022) 'Attention mechanism in neural networks: where it comes and where it goes', *Neural Comput & Applic*, 34, pp.13371-13385. Available at: <https://doi.org/10.1007/s00521-022-07366-3> (Accessed: 25/04/2025).
- Srinivasan, H. (2025) 'Beyond Common Reassurances of "It's OK"—The Reality of Anxiety in Autism', *Autism in Adulthood*. Available at: <https://doi.org/10.1177/25739581251366856> (Accessed: 13/06/2026).
- Steven, M. and Blakemore, C. (2004) 'Visual synaesthesia in the blind', *Perception.*, 33(7). Available at <https://doi.org/10.1068/p5160> (Accessed: 12/06/2026).
- Stevenson, R., Park, S., Cochran, C., McIntosh, L., Noel, J., Barense, M., Ferber, S. and Wallace, M. (2017) 'The associations between multisensory temporal processing and symptoms of schizophrenia.', *Schizophr Res.*, 179. Available at: <https://doi.org/10.1016/j.schres.2016.09.035>

Sutherland, S. (1998) 'Feature selection', *Nature*, 392. Available at: <https://doi.org/10.1038/32817> (Accessed: 04/06/2026).

Thomas, P., Zahorodny, W., Peng, B., Kim, S., Jani, N., Halperin, W., and Brimacombe, M. (2011) 'The association of autism diagnosis with socioeconomic status', *Autism*, 16(2), pp. 201-213. Available at: <https://doi.org/10.1177/1362361311413397> (Accessed: 10/06/2025).

Titchener, E. (1908) *Lectures on the Elementary Psychology of Feeling and Attention*. The Macmillan Company. Available at: <https://wellcomecollection.org/works/p9kqk88x> (Accessed:12/05/2025).

Treffert D. (2009) 'The savant syndrome: an extraordinary condition. A synopsis: past, present, future', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364(1522), pp.1351–1357. Available at: <https://doi.org/10.1098/rstb.2008.0326> (Accessed: 08/06/2026).

Trundle, G., Jones, K., Ropar, D., and Egan, V. (2022) 'Prevalence of Victimization in Autistic Individuals: A Systematic Review and Meta-Analysis', *Trauma Violence Abuse*., 24(4), pp.2282-2296. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10486169/> (Accessed: 06/06/2026).

Tsotsos, J., Itti, L., and Rees, G. (2005) 'A brief and selective history of attention', in L. Itti, G. Rees, and J. Tsotsos (eds.), *Neurobiology of Attention*. San Diego: Elsevier Academic Press, xxiii–xxxii. Available at: [10.1016/B978-012375731-9/50003-3](https://doi.org/10.1016/B978-012375731-9/50003-3) (Accessed: 06/07/2026).

Van de Cruys, S., de-Wit, L., Evers, K., Boets, B., & Wagemans, J. (2013) 'Weak priors versus overfitting of predictions in autism: Reply to Pellicano and Burr (TICS, 2012)'. *i-Perception*, 4(2), pp.95–97. Available at: <https://doi.org/10.1068/i0580ic> (Accessed: 06/07/2026).

Van der Hallen, R., Manning, C., Evers, K., and Wagemans, J. (2019) 'Global Motion Perception in Autism Spectrum Disorder: A Meta-Analysis', *J Autism Dev Disord*., 49(12), pp.4901-4918. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6841654/> (Accessed: 20/03/2026).

van Leeuwen, T., van Petersen, E., Burghoorn, F., Dingemanse, M., and van Lier, R. (2019) 'Autistic traits in synaesthesia: atypical sensory sensitivity and enhanced perception of details', *Philos Trans R Soc Lond B Biol Sci*., 374(1787). Available at: <https://doi.org/10.1098/rstb.2019.0024> (Accessed: 12/06/2026).

Verhoef, B. (2015) 'Fundamental Challenges for Autism Research: The Science–Practice Gap, Demarcating Autism and the Unsuccessful Search for the Neurobiological Basis of Autism.', *Medicine, Health Care and Philosophy*, 18, pp.443–47. Available at: <https://doi.org/10.1007/s11019-015-9636-7> (Accessed: 02/07/2026).

Vermeulen, P. (2012) *Autism as Context Blindness*. Kansas: AAPC Publishing.

- Volden, J. and Lord, C. (1991) 'Neologisms and idiosyncratic language in autistic speakers', *J Autism Dev Disord.*, 21 (2). Available at: <https://doi.org/10.1007/BF02284755> (Accessed: 12/07/2025).
- Volkmar, F. and Reichow, B. (2013) 'Autism in DSM-5: progress and challenges', *Molecular Autism*, 4(13). Available at: <https://link.springer.com/article/10.1186/2040-2392-4-13> (Accessed: 12/03/2025).
- Waltz, M. (2013) *Autism a Social and Medical History*. Hampshire: Palgrave Macmillan.
- Ward, E. and Scholl, B. (2014) 'Inattention blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness', *Psychon Bull Rev*, 22, pp. 722-727. Available at: <https://doi.org/10.3758/s13423-014-0745-8> (Accessed: 25/04/2025).
- Ward, J. and Banissy, M. (2015) 'Explaining mirror-touch synesthesia', *Cogn Neurosci.*, 6(2-3), pp.118-133. Available at: <https://doi.org/10.1080/17588928.2015.1042444> (Accessed: 14/06/2026).
- Ward, J., Hoadley, C., Hughes, J., Smith, P., Alison, C., Baron-Cohen, S., and Simner, J. (2017) 'Atypical sensory sensitivity as a shared feature between synaesthesia and autism', *Scientific Reports*, 7(41155). Available at: <https://doi.org/10.1038/srep41155> (Accessed: 12/06/2026).
- Watzl, S. (2017) *Structuring Mind: The Nature of Attention and how it Shapes Consciousness*. Oxford Academic. Available at: <https://doi-org.ezproxy1.lib.gla.ac.uk/10.1093/acprof:oso/9780199658428.001.0001> (Accessed: 21/01/2025).
- Watzl, S. (2023) 'What attention is. The priority structure account', *WIREs Cognitive Science*, 14(1). Available at: <https://doi.org/10.1002/wcs.1632> (Accessed: 17/01/2026).
- WebMD (2025) *What Is Hyperlexia?*. Available at: <https://www.webmd.com/children/what-is-hyperlexia> (Accessed: 18/05/2026).
- Weiskrantz, L. (2007) 'The Case of Blindsight', in M. Velmans and S. Schneider (eds.), *The Blackwell Companion to Consciousness*. Available at: <https://onlinelibrary.wiley.com/doi/book/10.1002/9780470751466> (Accessed: 10/03/2025).
- Wheeler, D., Jacobson, J., Paglieri, R., and Schwartz, A. (1993) 'An experimental assessment of facilitated communication', *Ment Retard.*, 31(1). Available at: <https://pubmed.ncbi.nlm.nih.gov/8441353/>
- White, A. (1964), *Attention*. Oxford: Basil Blackwell.
- Williams, D. (1996) *Autism: An Inside-Out Approach*. London: Jessica Kingsley Publishers.
- Williams, L., Light, G., Braff, D., and Ramachandran, V. (2010) 'Reduced Multisensory Integration in Patients with Schizophrenia on a Target Detection Task', *Neuropsychologia*, 48(10). Available at <https://doi.org/10.1016/j.neuropsychologia.2010.06.028> (Accessed: 12/06/2026).

- Williams, Z., and Gotham, K. (2022) 'Current and lifetime somatic symptom burden among transition-aged autistic young adults', *Autism research*, 15(4), pp.761–770. Available at: <https://doi.org/10.1002/aur.2671> (Accessed: 08/06/2026).
- Wilson, K. (2007) 'Does attention exist?', *British Journal of Undergraduate Philosophy*, 2(2), pp. 153-168. Available at: <https://www.research.ed.ac.uk/en/publications/does-attention-exist/> (Accessed: 07/06/2026).
- Wing, L. (1997) 'The History of Ideas on Autism: Legends, Myths and Reality', *Autism*, 1(1), pp. 13-23. Available at: <https://doi.org/10.1177/1362361397011004> (Accessed: 10/06/2025).
- Woods, S. and Estes, A. (2023) 'Toward a more comprehensive autism assessment: the survey of autistic strengths, skills, and interests', *Frontiers in psychiatry*, 14. Available at: <https://doi.org/10.3389/fpsyt.2023.1264516> (Accessed: 10/06/2026).
- Wu, W. (2014) *Attention*. Abingdon: Routledge
- Wu, W. (2019) *Review of Structuring Mind: The Nature of Attention and How It Shapes Consciousness*, by S. Watzl. *Mind*, 128(511), pp. 945–953. Available at: <https://www.jstor.org/stable/48595995> (Accessed: 10/02/2024).
- Yantis, S. and Jonides, J. (1990) 'Abrupt visual onsets and selective attention: Voluntary versus automatic allocation', *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), pp.121–134. Available at: <https://doi.org/10.1037/0096-1523.16.1.121> (Accessed: 20/06/2026).