



# MONASH University

# The emergence of autism in the Bayesian brain

Colin John Palmer

Bachelor of Behavioural Neuroscience (Hons)

*A thesis submitted for the degree of Doctor of Philosophy  
at Monash University in 2016*

Philosophy Department  
School of Philosophical, Historical and International Studies  
Faculty of Arts, Monash University

## **Copyright notice**

© Colin Palmer (2016). Except as provided in the Copyright Act 1968, this thesis may not be reproduced in any form without the written permission of the author.

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

# Contents

<b>Abstract</b> .....	1
<b>List of publications during enrolment</b> .....	3
<b>General declaration</b> .....	4
<b>Acknowledgements</b> .....	7
<b>Preface</b> .....	8
<b>Chapter 1 – Bayesian approaches to autism: Towards action and behaviour</b> .....	13
<b>Linking text between chapters 1 and 2</b> .....	55
<b>Chapter 2 – Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum</b> .....	60
<b>Linking text between chapters 2 and 3</b> .....	97
<b>Chapter 3 – Context sensitivity in action decreases along the autism spectrum: A predictive processing perspective</b> .....	99
<b>Linking text between chapters 3 and 4</b> .....	126
<b>Chapter 4 – Developing Bayesian accounts of autism: Evidence from statistical learning</b> .....	129
<b>Linking text between chapters 4 and 5</b> .....	146
<b>Chapter 5 – ‘Subtypes’ in the presentation of autistic traits in the general adult population</b> .....	149
<b>Linking text between chapters 5 and 6</b> .....	180
<b>Chapter 6 – The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism</b> .....	184
<b>Concluding remarks</b> .....	219
<b>References</b> .....	224
<b>Appendices</b> .....	259

## Abstract

The brain represents the state of the world around us based upon a stream of ambiguous sensory data. This requires the integration of different sources of information, including information received from different sense modalities, and information drawn from past experience. *Bayesian* theories of perception provide an approach to characterising how systematic differences in the probabilistic integration of information in the brain may underlie the differences that occur between individuals in their perceptual experience. The present thesis explores how perceptual, motor and social aspects of *autism spectrum disorder* (ASD) may emerge from variation in the neurocognitive processes described by Bayesian theories. A focus is on the *predictive processing* model of brain function, which links Bayesian theories of perception to neural mechanisms.

The thesis includes *empirical studies* that examine perceptual, sensorimotor and cognitive aspects of ASD. ASD and nonclinical autistic traits are first examined in the context of *body perception*; concerning, specifically, how the brain represents the state of the limbs based on visual, tactile and proprioceptive information, and how this sensory information is integrated with prior expectations regarding the body. These processes are investigated using a multisensory perceptual illusion: the *rubber hand illusion*. ASD is associated with a largely typical perceptual experience of this illusion, indicating intact multisensory integration in the context of body representation. A reduced influence of the illusion on motor function, however, supports a difference in ASD in the integration of expectations for limb position (influenced by the illusion) with conflicting sensory (proprioceptive) signals.

The integration of incoming information with existing expectations about the environment is also examined in the context of *statistical learning*. A computerised task is developed to assess the integration of iterative feedback with prior information as participants predict the location of a noisy set of visual markers. The hypothesis tested is that features of ASD are associated with a persistently higher weighting of incoming information in driving inference, at the expense of

expectations developed from recent experience; in this more cognitive context, these data provide evidence *against* this hypothesis, suggesting that a more context-dependent atypicality in information processing in ASD is more likely than that initially hypothesised.

The thesis also develops *theoretical* and philosophically-relevant treatments of how the symptoms of ASD may emerge from differences in brain function characterised within a Bayesian framework. This includes drawing on recent theoretical developments regarding the role of *volatility processing* and *action* in maintaining optimal inference on the environment. In addition, models of Bayesian inference in the brain can be extended to the social domain. For instance, the process of representing the mental states of other people (i.e., *theory of mind*) can be cast in terms of implicit inference on the external causes of sensory signals. In light of these expanded models of predictive processing, the differences in information processing hypothesised to occur in ASD have implications not just for perception, but also for motor behaviours, social cognition and social interaction.

The overall view that emerges is that diverse aspects of ASD may be captured in terms of how incoming sensory signals are integrated, in a probabilistic manner, with the brain's hierarchical and multimodal model of its environment. A promising direction for this field is in developing this idea in the context of action (i.e., active inference) and more recent models of how the brain estimates the optimal weighting of sensory information. This area of research has the potential to provide a nuanced perspective on the neurocognitive basis of ASD and the relationship between sensory mechanisms and autistic behaviours.

## List of publications during enrolment

Palmer, CJ, Paton, B, Hohwy, J, Enticott, P. (2013). Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum. *Neuropsychologia*, 51(10), 1942–1951.  
doi:10.1016/j.neuropsychologia.2013.06.020

Palmer, CJ, Paton, B, Enticott, P, Hohwy, J. (2015). 'Subtypes' in the presentation of autistic traits in the general adult population. *Journal of Autism and Developmental Disorders*, 45(5), 1291–1301. doi:10.1007/s10803-014-2289-1

Palmer, CJ, Paton, B, Kirkovski, M, Enticott, PG, Hohwy, J. (2015). Context sensitivity in action decreases along the autism spectrum: A predictive processing perspective. *Proceedings of the Royal Society B: Biological Sciences*, 282(1802). doi:10.1098/rspb.2014.1557

Palmer, CJ, Seth, A, Hohwy, J. (2015). The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism. *Consciousness & Cognition*, 36, 376–389. doi:10.1016/j.concog.2015.04.007

## **General declaration for thesis including published works**

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes four original papers published in peer-reviewed journals and two unpublished papers. The core theme of the thesis is to draw on Bayesian models of neurocognitive function to understand the mechanisms that underlie characteristics of autism spectrum disorder. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the candidate, working within the Philosophy Department under the supervision of Professor Jakob Hohwy, Associate Professor Peter Enticott, and Dr Bryan Paton.

The inclusion of co-authors for each thesis chapter reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research.

In the case of the chapters that comprise this thesis my contribution to the work involved the following:

<b>Thesis chapter</b>	<b>Publication title</b>	<b>Publication status</b>	<b>Nature and extent (%) of students contribution</b>
1	Bayesian approaches to autism: Towards action and behaviour	Not yet submitted	Developed ideas. Reviewed literature. Wrote the paper.  80%

2	Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum	Published	Contributed to the experiment design. Performed data collection and data analysis. Contributed to the interpretation of results. Wrote the paper.  50%
3	Context sensitivity in action decreases along the autism spectrum: A predictive processing perspective	Published	Contributed to the experiment design. Performed data analysis. Contributed to the interpretation of results. Wrote the paper.  50%
4	Developing Bayesian accounts of autism: Evidence from statistical learning	Not yet submitted	Contributed to the experiment design. Performed data analysis. Contributed to the interpretation of results. Wrote the paper.  Performed data collection and analysis for pilot experiments during development of the experimental paradigm.  60%
5	'Subtypes' in the presentation of autistic traits in the general adult population	Published	Contributed to the experiment design. Performed data analysis. Contributed to the interpretation of results. Wrote the paper.  70%

6	The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism	Published	Developed ideas. Reviewed literature. Wrote the paper.  80%
---	--	-----------	---

The content of each thesis chapter corresponds to the published papers, other than in minor details such as in the use of abbreviations. The numbering of sections within each chapter remains consistent with the published papers.

**Student signature:**

**Date:** 16/02/2016

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student and co-authors' contributions to this work.

**Main Supervisor signature:**

**Date:** 16/02/2016

## Acknowledgements

Thank you first and foremost to my primary supervisor, Jakob Hohwy. His deep engagement with this project and generosity with his time has helped to make it a body of work that I am proud of. He has also been a joy to work with: discussing ideas together was often the most fun part of the process. I have been privileged to work in the unique interdisciplinary environment that Jakob and Bryan set up, and on what, for me, has been a compelling and eye-opening area of research.

The work contained in this thesis has also benefited greatly from the involvement of my co-supervisors, Bryan Paton and Peter Enticott. Amongst other things, they have helped to maintain scientific rigour in the face of philosophical enthusiasm. I am also grateful for the involvement of other co-authors and colleagues, who are listed for each chapter. The reviewers of articles submitted for publication often surprised me with the level of their engagement with the work and the insight of their comments.

Lastly, I am grateful for my family and friends, who have supported me in many ways, not least in making it possible to keep a healthy perspective on the academic process. Dinah and Rob, Stephen and Clare, Catherine and Zac, Rosie and Oscar, Nick, and others...

## Preface

Autism spectrum disorder (ASD) is a neuropsychiatric developmental condition that is typically diagnosed at 3–4 years of age, and persists throughout life (American Psychiatric Association, 2013; Lai, Lombardo, & Baron-Cohen, 2014). Childhood is characterised by persistent deficits in social interaction, a strong preference for routine, repetitive motor behaviours, and sensory sensitivities. The precise presentation of these core diagnostic signs can differ considerably between cases. Associated difficulties with language, motor skills, and general intellectual functioning occur in some individuals who meet a diagnosis, but not others. In adulthood, social and sensory difficulties persist, contributing to poor educational, occupational, and relationship outcomes for the majority of those with ASD (Henninger & Taylor, 2013; Levy & Perry, 2011).

ASD was first defined on the basis of common patterns of distinctive behaviours exhibited in the child clinics of Hans Asperger and Leo Kanner in the 1940s. The conception of ASD has continued to evolve with time, to encompass a broader range of individuals who share similar characteristics, in recognition of sensory aspects of the condition, and with greater awareness of how challenges in childhood persist into adulthood (Silberman, 2015). Recent epidemiological studies estimate that worldwide prevalence is around 1% or greater (Baxter et al., 2015; Lai et al., 2014), with the most recent report of the Centres for Disease Control and Prevention estimating a prevalence of 1 in 45 in the US population (Zablotsky, Black, Maenner, Schieve, & Blumberg, 2015). ASD has a strong genetic basis, though one that is highly complex and heterogeneous (Geschwind & State, 2015), and diagnosis is made on the basis of behavioural observation alone. Critical insights into the physiological basis of ASD have emerged in part with the application of neuroimaging in recent decades (Ecker, Bookheimer, & Murphy, 2015); however, the link between neural atypicalities and the psychological symptoms of ASD are not yet well understood.

A recent trend in cognitive neuroscience is to draw upon *Bayesian probability theory* to characterise perceptual and cognitive mechanisms, and to aid in elucidating the physiological processes that realise these functions (e.g., Friston, 2005; Frith, 2007). The resulting picture of how the brain processes sensory information can be considered across a broad range of brain functions and, importantly, has promise in linking cognitive-, computational- and neural-level descriptions of psychological phenomena. This framework is therefore appealing for understanding how *differences* in our experience and interaction with the world can emerge, both across individuals and between groups. The present thesis explores how both social and non-social features of ASD can be understood within a Bayesian framework, and how cognitive and neuroscience research into ASD can be motivated and interpreted in this light.

*Chapter 1* provides a simple introduction to Bayesian theories of brain function, and reviews how systematic differences in the processes involved may account for important aspects of autistic perception. Of particular focus is *predictive processing*, which is a popular account of how Bayesian inference is implemented in the brain, and is important in part for tying descriptions of differences in information processing to predictions about brain function. *Chapter 1* also discusses lesser-explored aspects of Bayesian theories (including the role of action within predictive processing) that will be crucial for the field going forward in characterising the manner in which autistic symptoms might emerge from differences in neurocognitive function. While appearing at the beginning of the thesis, this chapter provides an overview of the current state of the field, and draws partly on work presented in the remaining components of the thesis.

*Chapters 2, 3, and 4* report empirical studies that investigate perceptual, motor and cognitive function in light of Bayesian theories of ASD. We first examine individual and group differences in sensory processing in the context of *body perception*, employing a perceptual illusion (the rubber hand illusion) that requires the brain to resolve between expectations about the body and conflicting multisensory inputs. *Chapter 2* reports evidence that autistic traits in non-clinical individuals modulate the perceptual and sensorimotor effects of this illusion, suggestive of

differences in information processing that are consistent with Bayesian theories of ASD. *Chapter 3* extends this experimental paradigm to a sample of adults diagnosed with ASD, and examines differences in the sensorimotor effects of the illusion in greater detail. Next, *Chapter 4* takes an alternative approach to testing Bayesian theories of ASD by examining how new information is integrated with prior experience in the context of statistical learning. Together, this research provides initial evidence for the differences in sensory processing hypothesised to occur in ASD, and highlights certain aspects of the Bayesian framework (e.g., the extension to action) that are useful for understanding and motivating empirical investigations.

A further component of the thesis relates to the conceptualisation of ASD as the high end of a *spectrum* of behavioural characteristics that encompasses the population at large, such that nonclinical individuals can be rated in terms of their level of *autistic traits*. Interestingly, there is evidence that cognitive and neurophysiological functions implicated in ASD vary in the general population with respect to autistic traits. Variation in the function of predictive processing mechanisms implicated in ASD may thus also contribute to individual differences in perception and behaviour in the general population. As noted, in this regard, the empirical work reported in *Chapters 2, 3, and 4* investigates perception and other functions partly in relation to nonclinical autistic traits. However, the manner in which autistic traits manifest in the general population is yet to be fully explored. In *Chapter 5* we examine in greater detail how autistic traits present in the general adult population with the use of cluster analysis and other related techniques.

Finally, the perceptual characteristics of ASD are a focus of much research in part because of how they may contribute to impaired *social* development. Bayesian theories of brain function provide a framework for relating social-cognitive functions and behaviours to aspects of sensory processing. For instance, our representation of others' *mental states* is a key paradigm in the field of social cognition; *Chapter 6* explores how this process can be understood in terms of probabilistic perceptual inference. A focus in this chapter is recent theoretical developments in Bayesian models of sensory processing; namely, the role of

*counterfactual sensorimotor predictions* in tying sensorimotor expectations to perceptual representations. These developments help to elucidate the demands on social cognition in an interactive context, with implications for how social differences may emerge in ASD from differences in predictive processing mechanisms.

The present document is a *thesis including published works*, for which each chapter was written to stand alone as a journal article. These chapters have not been rewritten for the purpose of the thesis. *Chapters 2, 3, 5 and 6* are published in peer-reviewed journals, and the published versions of these articles are included as appendices. *Chapters 1 and 4* are not yet published at the time of thesis submission. The thesis is of an interdisciplinary nature, relating most directly to neuroscience and psychology, while also touching on contemporary debates in the philosophy of mind. As noted, it includes both empirical studies and theoretical contributions. Brief linking text has been included between chapters to help the reader follow the broad connections between each piece. A short integrative discussion is included in the concluding remarks to the thesis.

## Declaration for Thesis Chapter 1

### Declaration by candidate

In the case of *Chapter 1*, the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
Developed ideas; reviewed literature; wrote the paper.	80%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Rebecca Lawson	Developed ideas; contributed to the writing.	N/A
Jakob Hohwy	Developed ideas; contributed to the writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date 16/02/2016
--------------------------	---	--------------------

Main Supervisor's Signature		Date 16/02/2016
-----------------------------------	---	--------------------

# Chapter 1

## Bayesian approaches to autism: Towards action and behaviour

Colin J. Palmer<sup>1</sup>, Rebecca Lawson<sup>2,3</sup>, Jakob Hohwy<sup>1</sup>

<sup>1</sup>*Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup>*Institute of Cognitive Neuroscience, University College London,  
London, WC1N 3AR, United Kingdom*

<sup>3</sup>*Wellcome Trust Centre for Neuroimaging, University College London,  
London, WC1N 3BG, United Kingdom*

## Abstract

Recently, our understanding of the sensory features of autism spectrum disorder (ASD) has been developed by drawing on *Bayesian probability theory* to describe systematic differences in the processing of sensory information in the brain. The key proposal is that ASD is characterised by a *greater weighting of sensory information* in updating probabilistic representations of the environment. Setting the appropriate weighting of sensory information is likely to be a complex neural process that relies upon inferences about the *volatility of the environment* and the interaction between *higher- and lower-level perceptual representations* in a hierarchical setting. This broader picture of inference in the brain will be important for understanding how the complexity of the autistic phenotype (e.g., its heterogeneity) emerges from atypicalities in Bayesian mechanisms. In addition, these considerations suggest that the stable trait in ASD may relate to finer mechanisms involved in the context-sensitive adjustment of sensory weighting rather than a persistent overweighting of sensory information *per se*. Moreover, in light of recent *sensorimotor* treatments of predictive processing (i.e., *active inference*), hypotheses regarding atypical sensory weighting in ASD have direct implications for the regulation of action and behaviour. Given that core features of ASD relate to how the individual *interacts* with the world around them (e.g., reduced social responding, repetitive behaviours, motor impairments and atypical visual sampling), the extension to action may be where Bayesian theories of ASD can yield the most critical insights into this condition.

## 1. Introduction

ASD is a common neurodevelopmental condition (approx. 1%) with a strong genetic basis, but lacks a clear neurological explanation tied to its cognitive profile (Lai et al., 2014). Heterogeneity in how symptoms express across the spectrum and the diverse range of social and non-social areas that tend to be affected within each individual contribute fundamental challenges in this regard (Happé, Ronald, & Plomin, 2006; Silberman, 2015). ASD lends itself to a *top-down* approach to cognitive neuroscience: we begin by identifying some key characteristics of cognition and behaviour, then move ‘downwards’ in an attempt to identify the underlying neurobiological mechanisms that would account for those characteristics. However, we can also consider a *bottom-up* approach. This approach begins by considering an overall theory of brain function, developed outside of the context of autism research, and then it moves ‘upwards’ to recover an account of how the autistic phenotype could emerge from within the constraints of that theory.

A general view of the brain’s function is that it works to *model* its environment, in this way ensuring that it can regulate its internal and external conditions for the sake of survival (Conant & Ashby, 1970; Friston, 2009). By drawing on computational theory regarding how models can be derived from sensory data, *predictive processing* has emerged as a general framework for understanding the functional organisation of the brain (Clark, 2013; Friston, 2005; Hohwy, 2013; Mumford, 1992). The theory is that top-down and bottom-up message passing across the cerebral cortex implements *hierarchical probabilistic inference* on the causes of sensory stimulation (Sections 2.1–2.2). It is currently a hot topic of research in cognitive neuroscience to explore the ramifications of this framework for a range of brain functions implicated in ASD. This includes sensory perception (Friston, 2005; Mumford, 1992), movement execution (Friston, Daunizeau, Kilner, & Kiebel, 2010; Shipp, Adams, & Friston, 2013), social cognition (Kilner, Friston, & Frith, 2007; Koster-Hale & Saxe, 2013) and autonomic bodily functions (Barrett & Simmons, 2015; Seth, 2013).

The emerging hypothesis for ASD is that incoming sensory signals are *weighted more highly* when integrated with the brain's existing model of the environment, such that perception (and other neural processes) are dictated to a greater extent by the present sensory data rather than prior or contextual information (Hohwy, 2013; Lawson, Rees, & Friston, 2014; Palmer, Paton, Hohwy, & Enticott, 2013; Pellicano & Burr, 2012; Van de Cruys et al., 2014) (Section 2.3). Setting the appropriate weighting of sensory information is a fundamental aspect of probabilistic inference, linked to modulatory neurotransmission in the brain, and made vulnerable to disorder (and individual differences) in part by the need to continually estimate the reliability of sensory information to maintain optimal inference. Many cognitive theories of ASD have been proposed in the 70 years since this condition was identified, with reoccurring themes across contemporary accounts including an imbalance between top-down and bottom-up processing of information in the brain (Happé & Frith, 2006; Mottron, Dawson, Soulères, Hubert, & Burack, 2006), the role of signal noise in sensory processing (Davis & Plaisted-Grant, 2015; Simmons et al., 2009) and distinctive learning or analytical styles (Baron-Cohen, 2009; Qian & Lipkin, 2011). It is enticing, therefore, to formalise these ideas within a general computational framework of brain function, with the prospect of tying cognitive-level descriptions more closely to their physiological implementation.<sup>1</sup>

A broad picture of the interacting mechanisms involved in regulating probabilistic inference in the brain will be necessary, however, to adequately explain how the complex clinical symptomology of ASD can emerge from Bayesian mechanisms. In particular, probabilistic inference in real-world environments necessitates *interactions between hierarchical and multimodal levels of environmental representation and estimates of environmental volatility* (i.e., the tendency for the causes of input to change over time) (Section 3.1–3.2). This latter demand is emphasised by recent computational work on probabilistic inference in changing environments by Christoph Mathys and colleagues (Mathys, Daunizeau, Friston, & Stephan, 2011; Mathys et al., 2014), which provides a method for modelling the

---

<sup>1</sup> See Brock (2014) for a perspective on how Bayesian accounts fit into the history of cognitive theories of autism.

role of volatility processing in perceptual inference. Differences in the finer mechanisms that underlie the context-dependent adjustment of sensory weightings may manifest in a more variable manner than a persistent overweighting of sensory information; this may be key to understanding the heterogeneity and variability in autistic characteristics (Section 3.3).

Moreover, ASD fundamentally reflects a difference in how the individual interacts with the world around them. Sensorimotor formulations of predictive processing (referred to as 'active inference'; Adams, Shipp, & Friston, 2013; Shipp et al., 2013) are therefore central to linking differences in inferential mechanisms to clinical symptoms (Section 4). In particular, the hypothesised differences in the weighting of sensory signals have direct implications for how the *balance* between perceptual updating and action is regulated, with important implications for the manner in which the nervous system interacts with the world to optimise its internal model. Under the predictive processing account, this determines where we sample from in the visual field, how we interact with others and the kind of motor behaviours that we engage in. Core social and non-social aspects of ASD can thus be cast in terms of differences in the regulation of active inference, including visual attention and visual search behaviours, social responsiveness, repetitive behaviours and differences in movement initiation. This extension to action and behaviour may be where Bayesian theories hold the greatest promise for illuminating the mechanisms underlying ASD.

## **2. Autism in a Bayesian framework**

### **2.1. An introductory picture of Bayesian inference**

Consider the task of determining the daily rainfall in your garden. The reading from your rain gauge will vary each morning even if the actual rainfall is constant, as strong winds and the like interfere with its accuracy. To get an estimate that is less susceptible to these unwanted variations, we can record the rainfall each morning and calculate the average of these measurements. In fact, we can update

our average after each new measurement, efficiently ensuring that our best estimate of rainfall is always available to inform our behaviour. This means that each new piece of information isn't interpreted alone, but is integrated into an existing model of rainfall. In this way, information from each measurement is added to the knowledge gained from all previous measurements.

The optimal manner of updating beliefs sequentially is with Bayesian inference. In this approach, beliefs are represented probabilistically. For example, our belief about daily rainfall might be represented as a distribution of probability values across a range of possible rainfall amounts (see *Figure 1a*). The process of Bayesian inference entails updating our existing belief (the *prior* distribution) with new information (the *likelihood* distribution) to form our new belief (the *posterior* distribution). We therefore need to represent three probability distributions to perform optimal inference; if we assume that these distributions are Gaussian, they can each be represented with just their mean ( $\mu$ ) and precision ( $\pi$ ).<sup>2</sup> The sequential updating of beliefs then proceeds as follows (Mathys et al., 2011), where  $x$  is the new measurement:

(1)

$$\mu_{\text{posterior}} = \mu_{\text{prior}} + \frac{\pi_{\text{likelihood}}}{\pi_{\text{posterior}}} (x - \mu_{\text{prior}})$$

where

(2)

$$\pi_{\text{posterior}} = \pi_{\text{prior}} + \pi_{\text{likelihood}}$$

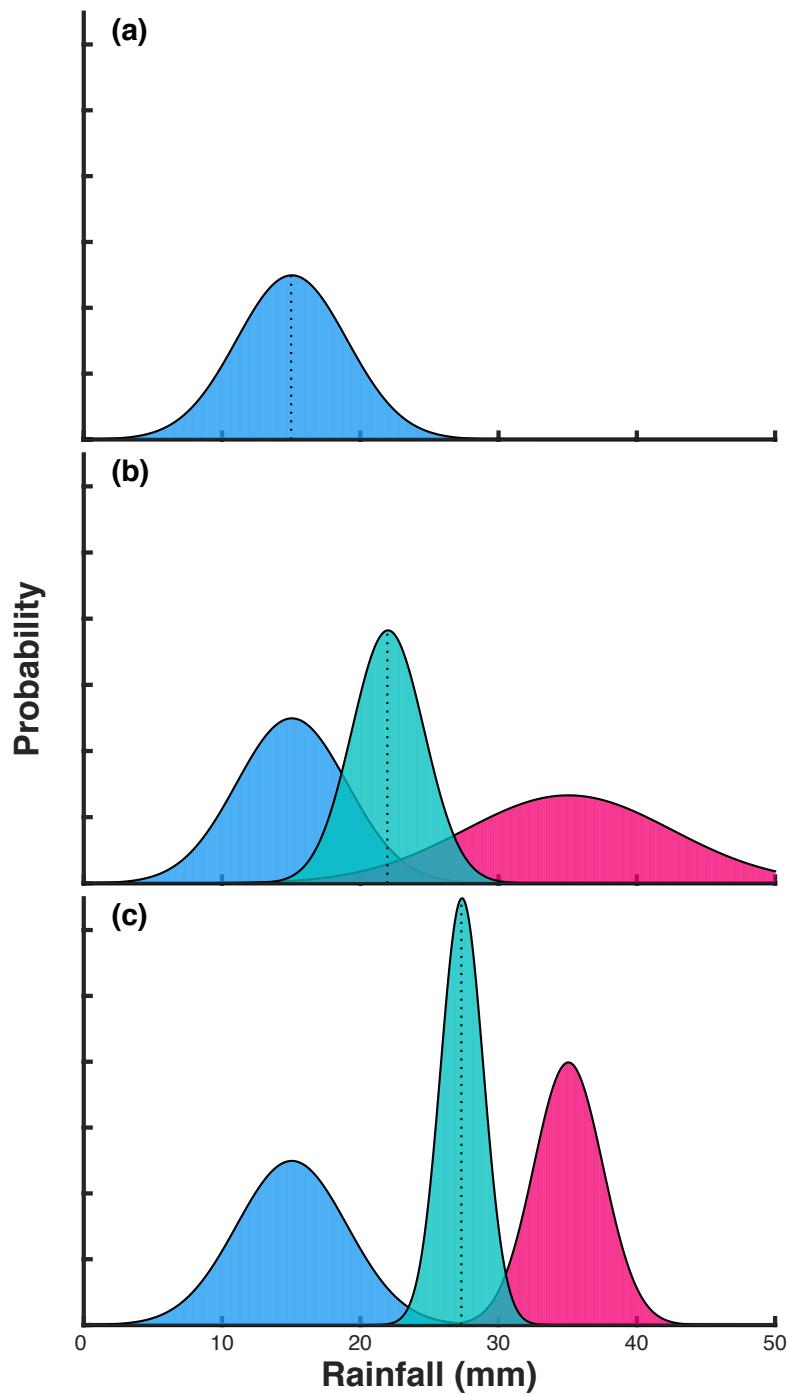
The mean of the prior belief ( $\mu_{\text{prior}}$ ) can be considered a *prediction* about what the new measurement will be; the particular reading on our rain gauge that we think is most likely to occur each morning. This is because the mean of a Gaussian

---

<sup>2</sup> Precision refers to the inverse of the variance.

probability distribution is also its maximum – the most probable outcome. The difference between this prediction and the measurement ( $x - \mu_{prior}$ ) is the *prediction error*. Hence, Bayesian inference can be done by iteratively updating predictions with the prediction error produced by each new measurement.

It is fundamental to probabilistic inference that the *extent* to which beliefs are updated in light of new information is set appropriately. The *precision* of the prior distribution indicates our confidence in our existing prediction, while the precision of the likelihood distribution represents the ambiguity inherent in the measurement (the noisiness of incoming data). Together, these two parameters give an indication of how reliable or informative prediction errors are expected to be regarding the true (hidden) state of the world. Prediction errors are therefore *weighted* by the estimated precision of the new information relative to the estimated precision of existing beliefs. *Figure 1b–c* illustrates the effect that this weighting term can have on what is inferred.



**Figure 1.** Bayesian inference.

(a) Belief about daily rainfall represented as a Gaussian probability distribution. The mean (dotted line) indicates the subjectively most probable rainfall.

(b) The posterior belief (green) is formed by integrating the existing belief (blue) with new information, where the latter is represented as a likelihood distribution (red). The mean of the likelihood is the newly collected data point ( $x$ ), while the uncertainty associated with this measurement is represented by the precision of the likelihood ( $\pi_{likelihood}$ ). The precision of a probability distribution reflects the degree of ambiguity in the state of the environment. For instance, a highly *precise* likelihood indicates that there are fewer states of the environment that have a high probability of causing the collected data; an *imprecise* likelihood means that there are many external states that could feasibly have caused the collected data.

(c) The extent to which the prior and the likelihood each influence the posterior depends upon their relative precision ( $\frac{\pi_{likelihood}}{\pi_{posterior}}$ ). In this example, the estimated precision of the newly collected data point is greater than in example (b), and as a result the inferred rainfall ( $\mu_{posterior}$ ; indicated by the dotted line) is based more closely on this new measurement.

The weighting term ( $\frac{\pi_{likelihood}}{\pi_{posterior}}$ ) is the *learning rate*. A high learning rate means that prediction errors drive inference about the state of the world to a greater extent. Conversely, a low learning rate means that prior information is afforded more weight in determining what is inferred. The learning rate ensures that beliefs are more highly sensitive to new measurements when we know little about the world (when our existing beliefs are imprecise), but less sensitive when we have already gathered plenty of information about the world (when our existing beliefs are highly precise). As we measure rainfall each morning, the precision of our belief increases, and as a result, the extent to which it is altered in response to new prediction errors is reduced. Similarly, when set appropriately, the learning rate ensures that a prediction error is more meaningful when we expect that the measurements we are taking are low in noise, such as if we live in a region where there are few unpredictable winds interfering with the reliability of our rain gauge.

Altogether, we can think of a prediction as a *hypothesis* about the *true underlying cause* of the collected measurements. If the data that we sample are drawn from a Gaussian distribution, then the hypothesis that will produce the least prediction error on average is the mean of this underlying distribution. In this way, minimizing prediction error over the long term, by updating predictions with precision-weighted prediction errors, is an effective way of representing the underlying cause of noisy data.

## 2.2. Bayesian inference in sensory cortex

From its seat within the skull, the brain faces a similar inferential challenge. What you see right now is driven by the firing rates of your retinal photoreceptors; however, for your brain to produce an accurate representation of the text in front of you, retinal signals must first be processed and made sense of by the sensory system. Significantly, the information that peripheral sense receptors provide is *ambiguous*, as there are many external states of the environment that can underlie a given pattern of sensory data. Similarly, perception is a *temporal* process, requiring the brain to continually adjust its representation of the environment while receiving an ongoing stream of sensory data. These considerations motivate the view of perception as a process of *unconscious inference*, requiring that hypotheses about the world are generated and continually updated to best account for sensory data (Gregory, 1980; Helmholtz, 1860; Kersten, Mamassian, & Yuille, 2004). In this view, ambiguity in incoming sensory signals is dealt with by drawing upon prior information regarding the relative probability of different possible states of the world, furnished by recent experience, early brain development and evolution. Given the key challenge of sensory ambiguity to accurate perception, it would not be surprising if evolution has shaped the brain's response to sensory stimulation to share features with optimal (probabilistic) inference.

As we have seen, probabilistic inference occurs when a system engages in iterative prediction-error updating. In the context of the brain, *predictions* are of sensory states and come to represent the hidden environmental causes of these states (edges, colours, objects and the like). *Prediction errors* communicate the

discrepancy between the actual sensory data and that predicted on the basis of the brain's current representation of the environment. *Precisions* indicate the estimated reliability of new sensory information and the confidence in existing predictions. The *predictive processing* theory of brain function is that reciprocal top-down and bottom-up message passing between sensory cortical areas mediates predictions and prediction errors, respectively, with their interaction regulated by the synaptic gain on cell populations that signal prediction-error, thus approximating optimal (precision-weighted) inference (Clark, 2013; Friston, 2005; Hohwy, 2013; Mumford, 1992). This theory has been developed in part to provide a computational explanation for various anatomical features of the cerebral cortex; for instance, that levels of processing can be defined hierarchically and that connections between cortical areas tend to occur reciprocally (Mumford, 1992; Shipp et al., 2013). There is much else to suggest that probabilistic concepts are useful for understanding the functional organization of the brain, including, for instance, psychophysical evidence that sensory estimates used in perceptual and motor systems are Bayes-optimal (see Vilares & Kording, 2011, for review) and neurophysiological evidence that stimulus-evoked responses in sensory cortex encode predictions and prediction errors (see Kok & de Lange, 2015, for review).

### **2.3. Autism as a variation from optimal inference**

Can a probabilistic outlook on brain function shed light on ASD? It is clear that biological differences in the message passing that constitutes predictive processing could lead to differences in perception and behaviour. Notably, optimal inference requires brain mechanisms that encode the present reliability of existing predictions relative to incoming sensory signals; in this way, prediction errors are given appropriate weighting in how the brain's representation of the environment is adjusted. In this regard, certain characteristics of ASD are suggestive of a *higher learning rate* in perceptual inference (Hohwy, 2013; Lawson et al., 2014; Palmer et al., 2013; Pellicano & Burr, 2012; van Boxtel & Lu, 2013b; Van de Cruys et al., 2014).

Pellicano and Burr (2012) argue that a *chronically reduced precision of (Bayesian) prior beliefs* in ASD underpins an array of sensory symptoms, psychophysical data and non-social behaviours. Van de Cruys et al. (2014) similarly suggest that a *high and inflexible weighting of prediction errors* can explain a wide range of social and non-social autistic characteristics. Brock (2012) notes that a tendency towards more sensory-driven Bayesian inference can be equally well explained by either *reduced* prior precision or *increased* likelihood precision (and Skewes, Jegindo, & Gebauer, 2014, devise an approach towards testing these alternative hypotheses in terms of signal detection theory). While this is an important conceptual point, each of these hypotheses reduces to a persistently high *learning rate* (weighting of prediction errors) and thus seemingly the same broad predictions regarding functional differences and physiological markers (i.e., atypical synaptic gain, described later in this section). We will now briefly summarise the evidence that motivates and supports these proposals.

A perceptual system that weights prediction errors more highly will tend to be more responsive to sensory stimulation and sensitive to finer fluctuations in sensory input. This fits well with hypersensitivities to sensory stimulation in ASD that are manifest in distress and sensory avoidance behaviours, first-person reports, and found across perceptual modalities in the psychophysical literature, including in certain tactile, auditory and visual discrimination thresholds and reduced habituation to repeated stimulation (e.g., Baron-Cohen, Ashwin, Ashwin, Tavassoli, & Chakrabarti, 2009; Blakemore et al., 2006; Puts, Wodka, Tommerdahl, Mostofsky, & Edden, 2014; Simmons et al., 2009). Unusual preoccupations with sensory stimulation (American Psychiatric Association, 2013; Zwaigenbaum et al., 2009) and enhanced attention to detail (or a more detail-oriented processing style), reported across a range of perceptual tasks that contrast processing of more local and more integrated perceptual elements (Happé & Frith, 2006; Mottron et al., 2006), are also suggestive of a persistently higher weighting of lower-level prediction error signals. Recent psychophysical evidence that adolescents with ASD show increased sensitivity to stimulus noise in motion perception similarly supports a greater sensitivity to fluctuations in the sensory signal from that predicted in sensory cortices (Zaidel, Goin-Kochel, & Angelaki, 2015).

Consistent with these aspects of perceptual experience and behaviour, adolescents with ASD show increased fMRI responses in primary sensory cortices to auditory, visual and tactile stimulation compared to typically developing controls, with the degree of activation related to behavioural symptoms of sensory over-responsivity (e.g., the tendency to become distressed in noisy environments; Green et al., 2015; Green et al., 2013). Evoked cortical responses to sensory stimuli also show increased variability across perceptual modalities (Dinstein et al., 2012; Haigh, Heeger, Dinstein, Minshew, & Behrmann, 2015) as well as differing dynamics compared to controls as the predictability of stimuli changes over repeated presentations (Gonzalez-Gadea et al., 2015; for review, see Van de Cruys et al., 2014 and Lawson et al., 2014). Most consistently, there is evidence that the reduction in neural response following repeated presentations of a stimulus (i.e., repetition suppression), a well-established phenomenon in neuroscience observed across sensory modalities, reflects increased predictability of the stimulus (Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008; Todorovic, van Ede, Maris, & de Lange, 2011); evidence that autistic traits are associated with reduced repetition suppression for both faces and non-social objects (Ewbank, Rhodes, et al., 2015; Kleinhans et al., 2009) is thus consistent with a failure to downregulate prediction error, as is a finding of reduced neural habituation to repeated tactile stimulation in adolescents with ASD (Green et al., 2015). (See Ewbank, von dem Hagen, Powell, Henson, & Calder, 2015, for a conflicting finding).

Moreover, an increased learning rate reduces the extent to which prior information informs what is perceived. The influence of past stimulation on current perception can be considered over various timescales. Sensory illusions occur when prior expectations that persist over *longer timescales*, reflecting statistical regularities in the world, bias perception away from a veridical representation of the present environment (demonstrated in Adams, Graf, & Ernst, 2004). Autistic individuals show reduced susceptibility to certain visual illusions, consistent with a higher learning rate (Mitchell, Mottron, Soulieres, & Ropar, 2010; conflicting evidence is reviewed in Section 3.3). In contrast, *perceptual adaptation* reflects an effect of *very recent sensory experience* on perception in the present

moment: exposure to a stimulus of one type biases perception of a subsequent (and immediately presented) stimulus away from that type. Adaptation effects are reduced in children with ASD for certain social and non-social visual stimuli (e.g., faces, Pellicano, Jeffery, Burr, & Rhodes, 2007; numerosity, Turi et al., 2015), consistent with a generally reduced influence of past experience on perceptual inference. Typical adaptation effects are found in ASD for certain other types of non-social stimuli (e.g., Karaminis et al., 2015), which Turi et al. (2015) suggest may reflect preserved adaptation for lower-level stimulus attributes.

Comparably, reduced *global processing* in ASD may reflect a reduced role for top-down predictions in *integrating* sensory features into a more broadly coherent or context-sensitive percept. Evidence for reduced global processing as a feature of autistic perception exists across a range of visual paradigms (Happé & Frith, 2006; Simmons et al., 2009), such as visuospatial object recognition (e.g., detection of embedded figures; Ring et al., 1999; Shah & Frith, 1983). Similarly, reduced global processing might be explained in terms of a persistent over-weighting of prediction error signals resulting in the brain's model of the environment being *overfitted* to noisy sensory data and thus failing to extract longer-term regularities in the sensory input.

In a somewhat more cognitive domain, Skewes et al. (2014) studied responses in a signal detection task that required statistical learning across trials in order to correctly categorise stimuli based on their visual characteristics. They found evidence that recent experience concerning the relative frequency of stimulus types influenced categorisation responses to a lesser extent in individuals higher in autism-like traits than those lower in these traits, consistent with a greater weighting of new information relative to prior information in ASD. Prior information is also essential for disambiguating how the observable behaviours of others reveal their hidden mental states such as intentions, emotions and beliefs (Kilner et al., 2007; Koster-Hale & Saxe, 2013), a capacity that may be selectively impaired in ASD, reflected in socio-cognitive developmental delays, reduced activity in neural regions typically recruited in social perception tasks, and reduced spontaneous mental-state attributions in adulthood (reviewed in Lai et al.,

2014). Thus, there is scope for illuminating the coexistence of sensory, cognitive and *social* deficits in ASD in terms of inference on sensory inputs (Hohwy & Palmer, 2014; Lawson et al., 2014; Palmer, Seth, & Hohwy, 2015; Van de Cruys et al., 2014). Van de Cruys et al. (2014) and Lawson et al. (2014) provide comprehensive summaries of the behavioural, psychophysical and neurophysiological data collected from autistic individuals that are consistent with chronic abnormalities in the weighting of prediction errors in inferential mechanisms in the brain.

In sum, a variety of autistic characteristics are suggestive of a greater weighting of sensory evidence in unconscious perceptual inference. In models of predictive processing in sensory cortex, the learning rate is set by the post-synaptic gain on distinct superficial pyramidal sub-populations that convey bottom-up prediction error signals between sensory cortical areas.<sup>3</sup> Synaptic gain is controlled in significant part by a number of interacting neuromodulatory systems, including those for which dysfunction is implicated in ASD (Feldman & Friston, 2010; Friston, 2009; Yu & Dayan, 2005) (see Lawson et al., 2014, for discussion in the context of ASD). Linking the perceptual and neurophysiological characteristics of ASD to Bayesian mechanisms has thus stimulated novel hypotheses regarding the neural, genetic and developmental basis of this condition (e.g., Lawson et al., 2014; Quattroki & Friston, 2014).<sup>4</sup>

More direct empirical investigation of these ideas are warranted, and will benefit in part from the ongoing effort to methodologically isolate the activity of neural populations that play distinct computational roles in predictive processing frameworks (Kok & de Lange, 2015). A clear and testable hypothesis is that cell

---

<sup>3</sup> Alternative hypotheses are suggested in the literature concerning the finer neural architecture of predictive processing (e.g., regarding the cortical layers that house neural subpopulations responsible for transmitting prediction error signals), reviewed in Kok and de Lange (2015).

<sup>4</sup> The process of optimising the weighting of sensory inputs across the brain during perceptual inference has been identified with *attention* in predictive processing theories (Feldman & Friston, 2010; Friston, 2009; Kok & de Lange, 2015). Hypotheses regarding atypical precision-weighting of sensory information in autism can thus also be cast in terms of atypical attentional mechanisms.

populations signalling prediction error exhibit enhanced gain in individuals with ASD (Van de Cruys et al., 2014). However, the interplay between mechanisms entailed by predictive inference complicate how simple neurocomputational atypicalities can be linked to heterogeneous behavioural symptomology. For this reason, a broader picture of probabilistic inference in the brain, portrayed in the remaining sections of this paper, will be crucial when using this theoretical approach to illuminate ASD and other neurocognitive conditions.

### **3. A broader picture of precision modulation**

Consider some different circumstances that may elicit prediction error in the brain (Mumford, 1992). First, prediction error may simply indicate that there is more to learn about the true value of a hidden state (e.g., the orientation of an edge in the visual field); in this case, prediction errors should be used to iteratively update the brain's estimate of the quantity that it is working to infer. Second, there may be irreducible noise in the brain's model of the environment, perhaps due to some genuine randomness in how the underlying cause generates sensory data over time; in this case, the existing predictions should be maintained even in the face of prediction error. Third, the predictions may be accurate, but still fail to fully explain the incoming data by itself; for example, it may be that the perceptual system correctly infers the presence of a cat, but that occluding branches and leaves alter the sensory signal, producing prediction errors that require the inference of additional states of the environment to explain. In this way, sensory uncertainty can indicate either *irreducible noise* or that there are *additional causes or interactions in the environment* that have yet to be modelled. The challenge here is that the 'one-level' treatment of Bayesian inference and predictive processing presented in Section 2.1 is inadequate to capture the demands on perceptual inference in real-world environments, as sensory input is always produced by a complex, interacting and fluctuating set of external causes. In the present and following sections we will delve into the mechanisms that this challenge entails, including the difficulty it poses to a simple Bayesian hypothesis of ASD, but also the promise that it holds for nuancing Bayesian accounts such that they are better

equipped to explain how the heterogeneous autistic phenotype emerges from systematic differences in sensory processing mechanisms.

### **3.1. Hierarchical representation**

We can begin to meet this type of challenge by considering the basic *hierarchical* and *multimodal* setting of predictive processing. Predictive inference in sensory systems is suggested to occur across a *hierarchy of inferred causes*, where higher levels engage in inference about causes that operate over greater spatial and temporal scales (Friston, 2005; Mumford, 1992; Shipp et al., 2013). This entails that the iterative prediction and prediction-error exchange that we saw in Section 2.1 occurs *recurrently* across the cerebral cortex, with top-down connections mediating predictions of sensory states at the level below, and prediction error computed at each level and signalled via feedforward connections to levels above. In the context of face perception, for instance, higher levels of the visual sensory hierarchy may represent global visual form and person identification, providing predictions regarding the activity at lower levels in the hierarchy that encode visual features that fluctuate over shorter timescales (e.g., perspective-dependent edges and colours). Similarly, hierarchical coding must occur partly in parallel across *different sensory modalities* (e.g., visual, auditory, tactile, interoceptive) and extend to multimodal representations of the environment that modulate lower-level predictions across these modalities. For example, inferences regarding speech content entail both visual and auditory predictions, and inferences regarding limb location entail both visual and proprioceptive predictions.

Representations of the environment thus occur at different levels of abstraction according to the hierarchical depth, with predictions generated at each level descending to act as a template or recreation of the expected pattern of sensory activity occurring at levels below (Mumford, 1992). In this way, hierarchical inference operationalizes the relationship between top-down and bottom-up sensory processing (and global and local perception) central to certain perceptual and cognitive characteristics of ASD (Happé & Frith, 2006; Mottron et al., 2006). A further feature of hierarchical inference is that it allows the different causes of

sensory input to be *deconvolved*; in other words, by modelling the interacting causal structure of the world internally, the brain can better account for patterns of sensory data produced (for example) by overlapping visual objects or simultaneous sound sources. Thus, the demands of perceptual inference in a real-world environment require that prediction error is minimised across extensive multimodal hierarchies of predictions. This complicates the picture of precision modulation in the brain, as it requires that the weighting of prediction errors is controlled at least part-independently in different modalities and at different levels of representation, this acting as a determinant on the relative influence of priors in different parts of the system (Friston, 2009). Minimising prediction error across the brain at a given point in time may entail variously *increased* and *decreased* weighting of prediction errors at different levels of the distributed hierachical system, such that the best explanation for sensory input emerges across the overall model.

The distinction between higher- and lower-level hypotheses is also important for capturing how prior experience modulates inference over the longer term in a Bayesian setting. For example, a ‘face’ hypothesis will entail a different set of priors to a ‘house’ hypothesis regarding lower-level visual features, reflected in both the value of predictions and the tendency for these low-level predictions to be adjusted in response to prediction error. For instance, the strikingly robust ‘hollow mask’ illusion occurs when the inside (concave) side of a facemask is perceived as a convex face – consistent with the statistical regularities of the world that we live in (Gregory, 1980). Research into the role of object type, orientation and lighting conditions in this paradigm indicates that the illusion reflects in part specific expectations about faces compared to other object types (Hill & Bruce, 1993, 1994). In predictive processing terms, this illusion can be explained as occurring when a high-level ‘face’ hypothesis entails predictions about the value of low-level visual features (related to convex object shape) and, moreover, reduces the learning rate at these lower hierarchical levels such that non-veridical predictions are maintained in the face of prediction error. This is equivalent to a highly precise Bayesian prior for convexity, but represented across a hierarchical predictive processing system, and, importantly, only invoked when higher-level face

hypotheses are selected. Thus, considering precision modulation in the setting of hierarchical inference is also important for operationalizing how expectations about the world that persist over longer timescales are brought to bear in a context-dependent manner.<sup>5</sup>

The nature of processing differences in ASD should be explored in the context of hierarchical processing. For instance, an account of ASD would ideally be able to specify the sensory modalities and level of hierarchical representation that differences in precision modulation occur, and, further, the secondary (or developmental) consequences of aberrant precision-weighting at one level of representation for the greater system. For example, there may be consequences for the development of higher-level representations if lower-level sensory cortex is characterised by an overweighting of prediction errors, and consequences for the development of multimodal representations if lower-level unimodal processing is characterised by atypical precision modulation.

### **3.2. Hierarchical inference in a changing world**

#### **3.2.1. Volatility in the environment**

The model of Bayesian inference that we have reviewed so far still carries a crucial limitation. When estimating daily rainfall in Section 2.1, we adjusted the weighting

---

<sup>5</sup> This consideration is important in the context of autism research, where certain atypical perceptual mechanisms, such as repetition suppression, have been implicated in social perception more so than for non-social objects (e.g., Ewbank, Rhodes, et al., 2015). In addition, the phenomenon of higher-level predictions being maintained in the face of prediction error can be demonstrated in the brain: using fMRI and population receptive field mapping, Kok and de Lange (2014) examined the *spatial profile* of neural activity in early visual cortex (V1) during viewing of illusory geometric shapes (Kanizsa triangles); neural activity (potentially reflecting prediction error) was variously increased or decreased in different parts of V1 depending on the region-specific consistency of the bottom up signal with the perceived shape. Specifically, if predictions for lower level features are drawn from inferences regarding higher-level object shape, predictions in this paradigm will correspond to illusory and non-illusory elements in different parts of the visual field, and thus will be associated with different degrees of prediction error when compared to the bottom-up signal.

of new information over time depending in part on how much we had already learnt. We can imagine that when seasons change, however, the information collected earlier in the year loses its relevance. Thus continuing to treat our prediction as highly precise would no longer give us an accurate inference about the (now different) daily rainfall. As the state of the environment changes, we need to adjust the learning rate to maintain optimal inference. Specifically, we may need to *increase* the learning rate when the underlying state of the world has changed, so that we learn about the new state. Thus, the common model of Bayesian inference (outlined in Section 2.1) only provides optimal inference in an *unchanging world*. In the real world, the causes of the brain's sensory input constantly change, meaning that a Bayesian model sensitive to the temporal demands of perception may be needed for a realistic picture of sensory processing.

Christoph Mathys and colleagues unpack this limitation by distinguishing between different forms of uncertainty (Mathys et al., 2014) (drawing on Payzan-LeNestour & Bossaerts, 2011; Yu & Dayan, 2005). Bayesian inference is sensitive to the idea that hidden states generate ambiguous sensory data – this is captured in the precision of the likelihood, and is termed *outcome uncertainty*. Bayesian inference also involves representing beliefs probabilistically, reflecting that the brain cannot ever be completely certain about what the hidden causes of its input are – this is captured in the precision of the posterior, and is termed *expected uncertainty*. Outcome uncertainty and expected uncertainty are reflected in how our inference about the state of the world changes as we sample new data, the aim being to update the hypothesis to be closer to the state of the world that generates this data. However, this is a different challenge to that of determining if the underlying cause of input has itself changed. This reflects *environmental uncertainty*, and is not captured in the common Bayesian model.

Changes to the environment can be both in the *causes* of sensory data and in the *noisiness* of sensory signals. In hand perception and movement, for instance, the brain must estimate the position of the hand based on visual and proprioceptive sensory signals. In addition to the position of the hand itself changing over time, the reliability of visual signals regarding hand location differ depending on task

context; for instance, visual information is more precise for the azimuthal location of the hand than for its depth relative to the eye (van Beers, Sittig, & Gon, 1998; van Beers, Wolpert, & Haggard, 2002). Similarly, the reliability of visual data may differ across visibility conditions (e.g., lighting, blurriness) even as the objects in our environment remain stable, with implications for how visual signals are integrated with other sources of information (Alais & Burr, 2004). To maintain optimal inference requires that the weighting of prediction errors be flexibly adjusted in response to contextual shifts in the reliability of the relevant sensory signals.

To complicate the picture further, we can make another distinction between two broad ways in which the world can change. In one case, the *parameters* of the hidden state change (i.e., the means and precisions; e.g., as the angle of the elbow joint changes from 30° to 45°, the hypothesised position of the arm ought be adjusted accordingly); in the other, the *quality* of the hidden state appears to change (e.g., as our eyes shift to foveate the house rather than the cat, the entire nature of what is causing our retinal firing changes). In the latter case, we recruit a new *model* to account for changes in the sensory data, rather than simply adjusting the parameters of an existing model. Therefore, to explain differences in perception between individuals, we may need to consider not only the iterative updating of parameters but also how the system switches between different broad, hierarchically-distributed hypotheses regarding the set of interacting causes that best explain its sensory input. This latter aspect of perceptual inference has yet to be explored in detail.

### 3.2.1. Predictive processing supplemented with volatility expectations

To accurately model the state of the environment over time, the brain requires expectations about *volatility* – how liable the environment is to change. The premise of perceptual inference, however, is that the brain has only indirect access to the real state of the environment; this includes environmental volatility, leading to a picture of predictive processing as being supplemented by *inferences about volatility*. The implication is that volatility expectations must play a role in

perception by modulating the learning rate used in inferences about the state of the environment (Mathys et al., 2011; Mathys et al., 2014). This would allow for an *adaptive learning rate* at each level of environmental representation, taking into account how the ideal weighting of prediction errors changes as we sample more data, as the underlying causes of input change, and as the reliability of sensory information fluctuates across environmental contexts. A hierarchical representation of the causal structure of the world is crucial to maintaining an adaptive learning rate, as causes that operate over longer timescales have implications regarding the expected volatility of causes inferred at lower levels, and for identifying fluctuations in the reliability of sensory signals across contexts.

Mathys and colleagues have developed a mathematical tool for modelling Bayesian inference modulated by volatility expectations: the *Hierarchical Gaussian Filter* (HGF; Mathys et al., 2011; Mathys et al., 2014). In this scheme, the confidence in predictions depends partly on inferred volatility estimates (and thus environmental uncertainty) in addition to expected uncertainty. This means that expectations about volatility modulate the denominator of the learning rate (e.g., in *equation 1* in Section 2.1), such that the weighting of prediction errors tends to decrease over time when the underlying environment is stable, then increase when a change in the underlying environment occurs. Moreover, volatility is represented *hierarchically*, such that inference is sensitive not only to the extent to which the world tends to change over time (volatility), but also to fluctuations in the level of volatility over time. In support of the idea that quantities of this kind are harnessed in the brain, neuroimaging data indicates that activity in distributed cortical and subcortical regions during an audio-visual learning task relate to precision-weighted prediction errors modelled by the HGF (Iglesias et al., 2013).

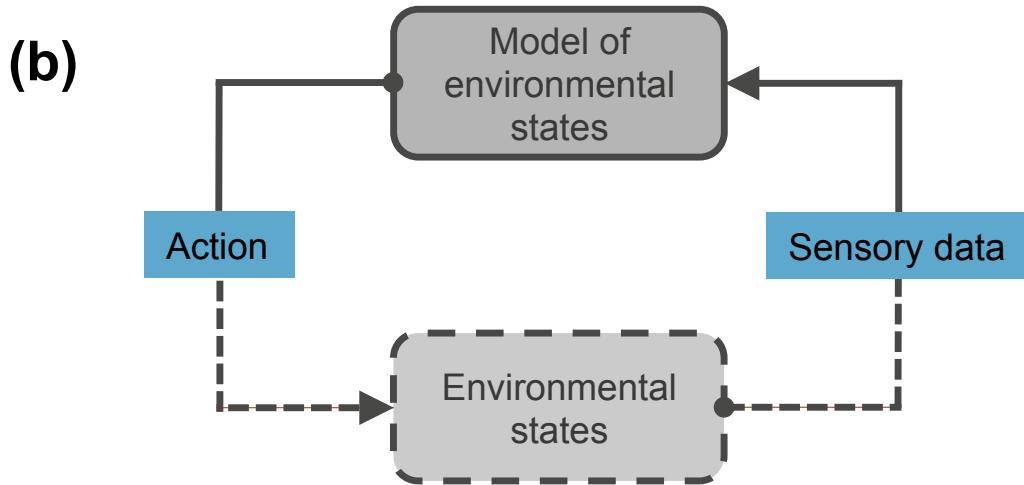
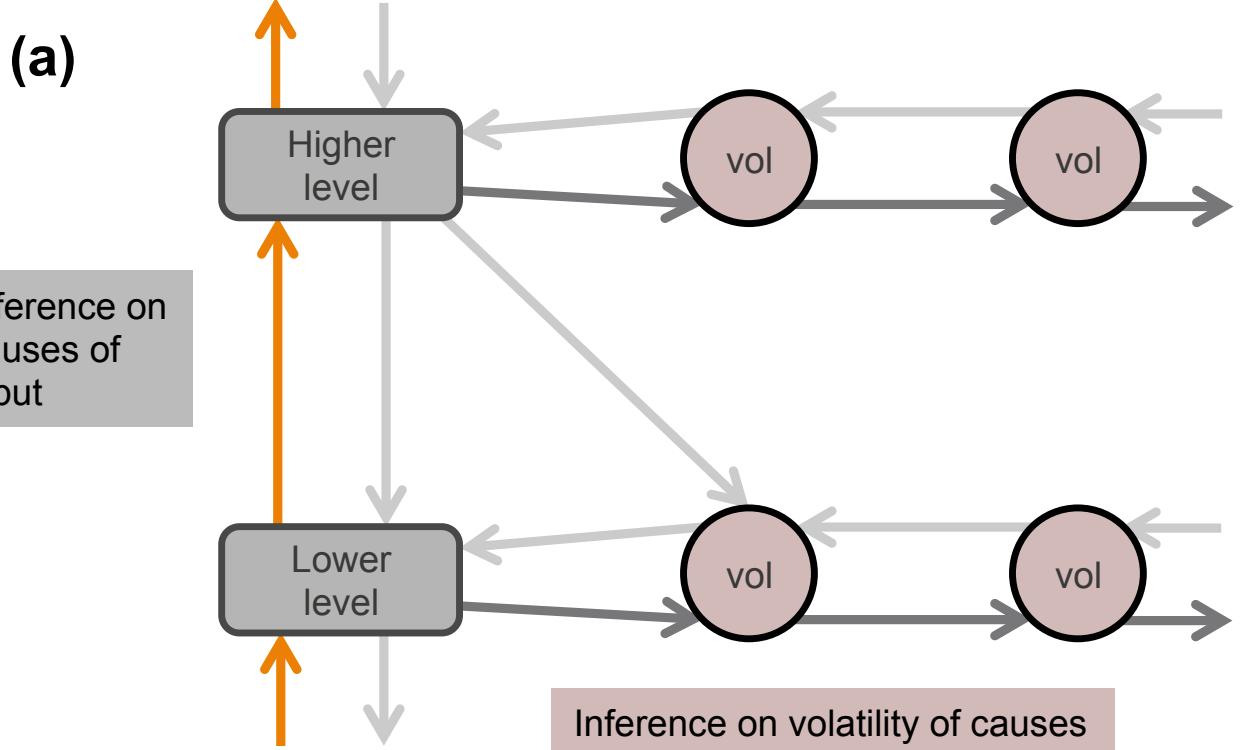
In the HGF, hierarchical inference of *volatility* is distinct from the *causal* hierarchy described in Section 3.1. In each class of hierarchy, estimates are adjusted via precision-weighted prediction errors, but in the causal hierarchy prediction errors relate to estimates of causes (e.g., the present location of a seen cat) rather than estimates of volatilities (how much the hidden location of the cat tends to change over time). We saw in Section 2.1 that when inferring the causes of sensory data,

predictions are of the *mean* of the hidden state. In contrast, the volatility hierarchy is one of *variances*; that is, the variance of the hidden state (indicating its tendency to change over time; its volatility), the variance of this variance (indicating how much the volatility changes over time), the variance of the volatility's variance, and so on.<sup>6</sup> In this respect, the causal hierarchy represents features of the environment that most clearly constitute the content of perception (see Hohwy, 2012; Hohwy, 2013; Seth, 2014), while volatility hierarchies modulate the processing that determines this inference by influencing learning rates at each level of the causal hierarchy.

Returning to our example of modelling rainfall, the data that we record each morning is determined by multiple interacting causes that operate over different timescales, each associated with differing degrees of volatility. For instance, cloud cover might change every day, while the effect that seasons have on rainfall remain relatively constant over the years. Moreover, various interactions might exist between volatility estimates and causal estimates at different levels; for instance, the inference of seasonal weather patterns that modulate sensory input over longer time scales may have implications for the expected volatility of inferences at shorter timescales. This amounts to a more convoluted picture of hierarchical inference in the brain, illustrated in *Figure 2a*, allowing for the highly context-dependent regulation of inference; and more closely capturing the nature of the environment that the brain is working to infer.

---

<sup>6</sup> Such a hierarchy might truncate after a small number of levels or keep going. Higher levels (of both causal and volatility hierarchies) correspond to longer timescales, however, meaning that eventually, as we go up the hierarchy, the information represented at higher levels will have little bearing on the shorter timescales that correspond to inference on the present state of the environment. In other words, for volatilities, higher levels will have a tendency to become uninformative.



**Figure 2.** Expanded models of predictive processing in the brain.

(a) Predictive processing models can be expanded to include inferences about volatility in a hierarchical setting (Mathys et al., 2011; Mathys et al., 2014). Bayesian inference on the external causes of sensory input occurs across a *causal hierarchy*, in which higher levels encode states of the environment that operate

over greater temporal and spatial scales, and include multimodal representations. At each level of the causal hierarchy, the weighting of prediction errors (learning rate) is modulated by a different set of hierarchical inferences concerning the *volatility* of the inferred causes. In both causal and volatility hierarchies, Bayesian inference occurs via the exchange of top-down predictions and bottom-up prediction errors (arrows in diagram). In the context of face perception, for example, a higher level in the causal hierarchy might correspond to inference on the global features of the face, sharing reciprocal connections with lower levels that model finer details of the face, and each of these levels in the causal hierarchy may be associated with volatility estimates provided by corresponding volatility hierarchies. Various interactions might exist between volatility estimates and causal estimates at different levels; for instance, the inference of certain higher-level causes that modulate sensory input over longer time scales may have implications for the expected volatility at lower levels of the causal hierarchy.

*(b) Active inference* extends predictive processing models to include action (Friston et al., 2010). The sensory data that the brain receives is generated by (hidden) external states of the physical world, which includes the states of the body. The brain models these external states internally based upon the sensory data (and prior expectations, expectations about volatility, etc.). Meanwhile, the brain can modulate the external states of the world via action, which influences the sensory data that is received. Minimising prediction error over time entails both *revising the model* to better account for sensory data, and *acting* to control the sensory data that is received.

### **3.3. Is a broader picture of precision modulation necessary to capture the characteristics of autism?**

In Section 2, we saw that several characteristics of ASD were suggestive of a higher learning rate in perceptual inference. This underlies proposals of *chronic* differences in precision weighting in ASD, framed within the general context of Bayesian inference (Pellicano & Burr, 2012) and within the predictive processing theory of inference in the brain (Van de Cruys et al., 2014). In Sections 3.1–3.2, we explored challenges to perceptual inference that require the weighting of prediction errors to be highly context-dependent, controlled in part by volatility expectations and the interaction between higher- and lower-level representations of the environment. This raises the possibility that the underlying trait in ASD is not a persistently higher weighting of prediction errors *per se*, but rather a difference in the mechanisms that control the context-sensitive adjustment of precisions. For instance, a stable difference in volatility processing might manifest as unusually high weighting of prediction errors in some contexts but not others. In the present section, we will consider several reasons why these underlying mechanisms implicated in the control of precision weighting in the brain may play a necessary explanatory role in Bayesian accounts of ASD.

#### 3.3.1. Accounting for empirical data

An initial reason to investigate beyond the possibility of chronic differences in precision weighting is that such accounts make *strong predictions* about perceptual and cognitive performance that are not always apparent in ASD.

One example is in the case of susceptibility to perceptual illusions, which is a key paradigm for examining how (implicit) expectations about the world modulate perceptual experience. While there is convincing evidence that susceptibility is reduced in ASD for some visual illusions, in particular the Shepard illusion, which relies on depth cues regarding the shape of a familiar object (Mitchell et al., 2010) (see also Bölte, Holtmann, Poustka, Scheurich, & Schmidt, 2007; Happé, 1996), it is also the case that individuals with ASD experience some perceptual situations that

draw upon prior knowledge very similarly to controls. Certain visual illusions seem to be experienced as usual in ASD, including the Muller-Lyer and Titchener illusions (Hoy, Hatton, & Hare, 2004; Ropar & Mitchell, 1999, 2001) (see Mitchell & Ropar, 2004, for review). Similarly, autistic traits in the general population negatively correlate with susceptibility to some, but not other, visual illusions (Walter, Dassonville, & Bochsler, 2009). Another example is the multisensory *rubber-hand illusion*, in which synchronous stroking of the subject's (hidden) limb together with a (visible) prosthetic limb elicits an illusion of feeling touch on the prosthetic limb (visual-tactile integration) and a sense that the prosthetic limb is a part of one's body (Botvinick & Cohen, 1998; Ehrsson, 2012). Susceptibility to this illusion is modulated by prior expectations about the body; for instance, there is evidence that the illusion is constrained by implicit expectations about the visual appearance of the body (Tsakiris & Haggard, 2005) but that this constraint can be overcome when visual-tactile experience immediately prior to the illusion favours unusual bodily representations (Hohwy & Paton, 2010). Adults with ASD exhibit the core perceptual effects of this illusion as robustly as controls, though show reduced effects of the illusion on subsequent reaching movements (Palmer, Paton, Kirkovski, Enticott, & Hohwy, 2015) (see also Cascio, Foss-Feig, Burnette, Heacock, & Cosby, 2012; Palmer et al., 2013; Paton, Hohwy, & Enticott, 2012), suggesting that the influence of priors on perception is not chronically diminished in ASD, but rather shows more subtle and context-dependent effects.

In the more cognitive domain of statistical learning, there is evidence that autistic traits in the general population are unrelated to chronic differences in the weighting of prediction errors (*Chapter 4*). This evidence comes from a study that examined trial-by-trial learning in a task that requires explicitly predicting the location of a visual marker based on statistical regularities. Similarly, Robic et al. (2015) examined trial-by-trial learning over the course of a decision-making task in which the (hidden) probabilities of reward associated with choice options were either *stable* over a block of trials or *fluctuated* during the course of a block (i.e., constituting a type of volatility). Adults with ASD had particular difficulty in performing the task in the volatile context, though performed similarly to controls when the task contingencies were stable over time. This suggests that these

participants were able to learn how to perform well when the outcomes were uncertain, but less able to deal successfully with volatility (i.e., when the underlying contingencies change). While this task examines a cognitive rather than sensory function, it constitutes initial evidence that volatility processing is impaired in ASD rather than the integration of prior and new information *per se*. The anterior cingulate cortex has been implicated in volatility learning in decision-making tasks of this nature (Behrens, Woolrich, Walton, & Rushworth, 2007), and atypical functioning in this region has been discovered in ASD (e.g., Di Martino et al., 2009; Dichter, Felder, & Bodfish, 2009; Thakkar et al., 2008).

### 3.3.2. Addressing conceptual challenges to neurocognitive accounts of autism

Furthermore, several challenges central to an adequate neurocognitive account of ASD are better addressed within the broader picture of hierarchical inference. First, *heterogeneity* is a defining feature of ASD (Newschaffer, Fallin, & Lee, 2002). For instance, individuals with ASD variously report (or exhibit) *hyper-* or *hyporesponsivity* to sensory stimulation in multiple sensory domains, reflected partly in either sensory-seeking or sensory-aversive behaviours (e.g., lower thresholds to touch detection, but either aversion or attraction to flickering lights; Baranek, David, Poe, Stone, & Watson, 2006; Marco, Hinkley, Hill, & Nagarajan, 2011; Rogers & Ozonoff, 2005). Similarly, high variability in autistic samples is reported across a number of visual psychophysical paradigms (Simmons et al., 2009). Individual social tendencies can entail marked detachment from social interaction with others or only more subtle atypicalities in conversational and other interactive behaviours. Intellectual functioning can similarly vary from that of severe disability to no disability at all. Neurocognitive explanations of ASD, which try to extract some degree of commonality from the assortment of individuals that fall under the ASD diagnosis, are thus more plausible if amenable to how heterogeneity arises through the hypothesised deficit. While a (chronically) high learning rate captures important aspects of autistic perception, a more nuanced picture of the context-dependent modulation of the learning rate, dependent on volatility expectations, higher-level expectations regarding lower-level sensory uncertainty, and different swathes of the hierarchy that may be impaired, is better

equipped to address how a core deficit (in precision modulation) plays out differently across individuals.

Accounting for the *diversity* of symptoms in ASD is a similar explanatory challenge. Social interaction and communication are central to defining ASD, but exist alongside non-social symptoms that span perceptual, motor, behavioural, and cognitive domains. The most widely influential neurocognitive theories of ASD each address certain symptoms much more directly than others; for example, the theory of mind deficit hypothesis (Baron-Cohen, 1997) addresses difficulties in understanding the behaviours of others, while the weak central coherence and enhanced perceptual functioning theories (Happé & Frith, 2006; Mottron et al., 2006) are based on perceptual and non-social cognitive differences in ASD. This leads to a common question of how social symptoms engender non-social deficits (e.g., by affecting learning in early development), or, alternatively, how general differences in sensory processing in the brain extend to the social domain. The approach that can be taken for Bayesian accounts is to *equate* important aspects of social and non-social processing: specifically, our representation of others' mental states can be cast in terms of Bayesian inference on the hidden external causes of sensory data, where this sensory data relates to the behaviours of others (Hohwy & Palmer, 2014; Kilner et al., 2007; Koster-Hale & Saxe, 2013; Palmer, Seth, et al., 2015). Similarly, we will see in Section 4 that differences in non-social sensory behaviours and social interaction may each reflect how the balance between action and perception is maintained in an inferential system. To make such an account plausible, an adequate framework is required for distinguishing between social and non-social inference and for defining their interaction; for instance, to explain why social functions are particularly vulnerable to general differences in sensory processing in the brain (Happé et al., 2006). To meet this challenge within Bayesian theories, we need to appeal to the broader hierarchical setting of perceptual inference (discussed in Hohwy & Palmer, 2014; Palmer, Seth, et al., 2015). For instance, commonalities (e.g., automaticity) between perception and implicit social-cognitive processes can be examined partly in terms of the *hierarchical depth* of inferred causes. The notion that social cognition is particularly vulnerable to differences in information processing (compared to

other perceptual or cognitive processes) due to the complexity of social situations can similarly be operationalized to an extent in a Bayesian framework; for example, inferring others' mental states may require modelling the influence of relatively long term causes in the environment on sensory input and a greater dependence on the context-dependent adjustment of precision expectations, and social situations entail a unique requirement to model the influence of our own mental states on those of others (Hohwy & Palmer, 2014; Palmer, Seth, et al., 2015).<sup>7</sup>

Finally, core symptoms of several psychiatric conditions other than ASD have been similarly characterised in terms of inferential mechanisms, including schizophrenia and Parkinson's disease (e.g., Edwards, Adams, Brown, Parees, & Friston, 2012; Fletcher & Frith, 2009). Notably, positive and negative symptoms of schizophrenia are suggested to be indicative of reduced precision of prior beliefs within a hierarchical predictive processing setting (Adams, Perrinet, & Friston, 2012; Adams, Stephan, Brown, Frith, & Friston, 2013). Accounting for ASD in terms of a chronically altered learning rate is therefore likely only to be the beginning of an adequate explanation for this condition; rather, a more nuanced description is needed of aberrant precision-modulation in the autistic brain to account for how ASD differs from other psychiatric conditions, taking into account developmental trajectories, the location in the brain (or sensory hierarchy) that primary and secondary differences occur, and mechanisms underlying precision modulation (e.g., the role of volatility expectations or higher-level contextual modulation). A recent example in this regard comes from Quattrocki and Friston (2014), which

---

<sup>7</sup> There is compelling evidence from large-scale studies that the behavioural traits that define the key DSM-IV diagnostic domains of autism (relating to social interaction, communication, and restricted behaviours) vary *independently* across the general population to a significant extent, meaning that these broad symptom domains may have separate genetic and neurocognitive causes despite their co-occurrence being necessary for a diagnosis of autism (reviewed in Happé et al., 2006). This argues in part against an attempt to explain social and non-social features of autism with appeal to the same underlying Bayesian mechanisms. However, these studies do indicate a tendency for impairments in DSM-IV domains to co-occur to a moderate extent and for significant overlap in their genetic causes. Furthermore, more specific cognitive, motor and perceptual characteristics implicated in autism may to some extent interrelate differently than the broad behavioural domains used in diagnostic procedures.

reviews evidence that ASD may stem from differences in the precision weighting of *interoceptive* sensory signals in early childhood, controlled by oxytocinergic neuromodulation. Predictive processing in the context of interoceptive sensory signals, and their hierarchical interaction with exteroceptive signals, has been linked to the experience of emotion and aspects of self-experience (such as bodily ownership) as well as bodily homeostasis (Barrett & Simmons, 2015; Seth, 2013). Impairment in the ability to modulate the salience of interoceptive sensory signals context-dependently (e.g., in response to relevant visual cues) is correspondingly suggested by Quattroki and Friston (2014) to compromise the generation of internal models of the self, with implications for the typical development of various aspects of social functioning, including observational learning, theory of mind and emotional responding (see Brewer, Happe, Cook, & Bird, 2015, for commentary).

## 4. Precision modulation and behaviour

### 4.1 Prediction-error minimisation and action

In the predictive processing model, perception occurs as prediction error in the sensory system is minimised by revising (hierarchical) cortical predictions in light of incoming sensory signals. In this way, the brain's representation of the world changes to better account for sensory data; experientially, the perceived state of the world changes. An alternative way in which prediction error can be reduced is by *acting on the world* to bring sensory input closer in line with that predicted (Friston et al., 2010). Consider the case of proprioception: descending predictions regarding proprioceptive sensory signals can be made to better match incoming data by either *revising* these predictions or *moving* to change the position of the body to that predicted. Thus, we can make a distinction between *perceptual inference* and *active inference* in their direction of fit in reducing prediction error (Friston et al., 2010; this distinction in terminology is provisional, see Section 4.3). *Figure 2b* illustrates the resultant situation of a biological agent schematically: sensory states are determined by (hidden) external states of the physical world,

which the agent models internally; meanwhile, the organism can act on the external states of the world, modulating its own sensory input. In this way, the agent is an active part of the causal structure of the world that generates the sensory data that its sensory system is working to predict.

These considerations furnish a unique perspective on movement execution and the functional organisation of the motor system (Adams, Shipp, et al., 2013; Friston et al., 2010; Shipp et al., 2013). Specifically, descending projections from motor cortex can be regarded as mediating *predictions* of proprioceptive states at the level of the spinal cord (rather than *motor commands*, as these signals are more typically conceived). Prediction error occurs at peripheral levels when proprioceptive states fail to match corticospinal predictions. Unlike in the perceptual system, however, peripheral prediction errors are resolved via spinal reflexes that engage the muscles to bring proprioceptive input in line with that predicted. In this way, motor predictions act as a set point, indicating the desired (or *expected*) bodily state that peripheral reflexes work to maintain. Hence, movement is initiated when (initially) *inaccurate* beliefs about proprioceptive states are realised by engaging spinal reflexes. This is made possible by the close relationship between muscular activity and the stimulation of proprioceptive sense receptors; for example, the firing of (sensory) muscle spindles is directly controlled by muscular contractions that alter the length of the muscles.

Descending motor signals (i.e., proprioceptive predictions) in the predictive processing account are the result of a multimodal, hierarchical model of the world, where this model includes a representation of the agent itself. Therefore, *exteroceptive* prediction errors might also be resolved through action; for instance, if visual predictions modulate proprioceptive predictions, the latter may induce movement that shifts the eyes to bring visual input in line with that predicted. In this way, active inference can be considered more broadly than proprioceptive control. An important example in this regard is that *bodily homeostasis* can be conceptualised as active inference on *interoceptive* sensory signals (Seth, 2013): falling blood glucose levels, for example, can be met by either *revising* interoceptive predictions (and failing to survive if the actual glucose concentration

falls below a safe limit) or by *acting* to maintain interoceptive signals within an expected range (via autonomic reflexes or motor behaviour). In some cases, interoceptive expectations ought to shape proprioceptive and exteroceptive predictions to drive behaviour (for instance, food seeking) that results in obtaining the predicted interoceptive states. This highlights the importance of a multimodal hierarchical system in which exteroceptive, proprioceptive and interoceptive expectations interact over different time scales to result in adaptive behaviour.

#### **4.2 Varieties of active inference**

Action as discussed so far can be understood as *confirmatory* active inference, in which the nervous system acts to make its beliefs about the environment true. Predictions for multimodal sensory activity are shaped by evolution, early development and recent experience, and function to keep the organism within a limited set of states favourable to survival; in this way, concepts of utility and adaptation can be captured in terms of hierarchical predictions that are constrained by natural selection and brain development (Friston, 2010).

In their classic formulations of unconscious inference, Hermann Von Helmholtz and later, Richard Gregory (Gregory, 1980; Helmholtz, 1860), emphasised the *investigative* role for action in testing the brain's beliefs about the world; for instance, the hypothesis that a table is a current cause of visual stimulation can be evidenced by moving about the room to sample consequent changes in the visual stream. In the social domain too, conversation is a richly interactive process, in which our speech, facial expressions and other behaviours interact with others' intentions, emotions and beliefs to garner us evidence about these internal psychological states (discussed in Palmer, Seth, et al., 2015). In this way, action plays a further role in inference by *reducing uncertainty* in the state of the environment. This can be distinguished as *disambiguatory* active inference, in which action helps to accumulate evidence for an existing hypothesis or resolve between competing hypotheses (delineated further in Seth, 2015a). This in part captures the benefit of *exploration* to an inferential system, even if it may come at the cost of transiently increased prediction error. In practice, we are always

interacting with the world, through saccades, speech and more overt movements; and as a process of disambiguatory sampling, these movements and behaviours are deeply ingrained in perceptual inference, with prediction error minimisation over the longer term entailing a continuing synthesis of action and perceptual updating.

To this end, active inference suggests that there would be benefit in deploying *counterfactual* predictions; that is, predictions of the sensory consequences of *potential* actions (Friston, Adams, Perrinet, & Breakspear, 2012; Seth, 2014, 2015a, 2015b). This is distinct from predictive processing as typically discussed in that it involves encoding a range of *possible future causes* of sensory input and the associated precisions expected (contingent upon actions that could be made), rather than modelling the present state of the environment. The importance of counterfactual predictions is in allowing for *action selection* on the basis of how alternatives for action are expected to reduce uncertainty in the brain's representation of the environment. The implication is that behavioural patterns (including visual search strategies, haptic exploration, and interactive social behaviours, for example) are driven in part to sample the world such that perceptual representations are made more precise. In the context of oculomotor control, Friston et al. (2012) demonstrate that saccadic visual search strategies can be modelled in terms of Bayesian inference driven by uncertainty-reduction. Beliefs of this type can be represented in terms of a salience map that indicates from where (in the visual field, for example) the system expects sampling will maximise its confidence in predictions about the world; for instance, an area of high salience might be the eyes of a face, the foveation of which will maximally disambiguate the type of object that we are looking at (or the mental state of the person we are interacting with). In this way, the inferential system *selectively samples* its environment through active inference to optimise its internal probabilistic model.

### 4.3 Balancing action and perception with precision modulation

The possibility of resolving prediction error through either perception or action poses a fundamental question: what mechanism determines *when* we act and when we update predictions instead? What prevents the brain from failing to act at all, instead adjusting its predictions to resolve prediction error? To initiate action rather than perceptual updating, the process required is *sensory attenuation*: the downregulation of sensory prediction errors such that predictions that are a poorer fit for the current input are nevertheless able to predominate and drive action (Brown, Adams, Parees, Edwards, & Friston, 2013; Brown, Friston, & Bestmann, 2011). This mechanism is important for understanding the implications of the hypothesised differences in precision weighting in ASD for the emergence of clinical symptomology.

Picture two hypotheses that your brain may have regarding the position of your hand. One hypothesis is that your hand is stationary by your side. An alternate hypothesis is that your hand is out in front of you, gripping the coffee cup. These hypotheses each translate to distinct patterns of proprioceptive signals for perception and action; that is, different descending (hierarchical) predictions in proprioceptive sensory and motor systems in the nervous system. If your arm really *is* at your side at present, then the first hypothesis will minimise sensory prediction error more so than the other. By attenuating the feedforward signal in sensory cortex, however, the alternate hypothesis can persist in the face of prediction error. As this hypothesis furnishes *motor* predictions that fail to match the actual proprioceptive signals, spinal reflexes will be engaged to extinguish motor prediction error, bringing your arm to the coffee cup.<sup>8</sup>

---

<sup>8</sup> Attenuation of the (predictable) sensory consequences of self-generated movements has received much empirical attention, and is discussed in the context of motor control theory (e.g., Blakemore, Frith, & Wolpert, 1999); in this model, predictions of the sensory consequences of self-generated actions are compared to the actual sensory consequences, with the degree of mismatch related to how strongly the sensory consequences of the action are perceived. Sensory attenuation in the context of active inference differs, however, in that bottom-up sensory signals are broadly attenuated (i.e., it is not just those signals that are predicted that are attenuated), such that sensory evidence that conflicts with the

Importantly, sensory attenuation is simply optimal precision weighting as we have already seen it, and thus relies upon mechanisms related to volatility expectations and hierarchical interactions as described in Section 3. As noted in Section 2, Bayesian inference isn't limited to simply minimising *immediate* prediction errors, but aims rather to infer the *underlying* cause of noisy data and thus minimise prediction error over the longer term. This requires that the precision of the prediction and the precision of sensory evidence are constantly estimated. When precision estimates tip in the direction of a *lower* weighting of prediction errors, predictions will be maintained even in the face of conflicting sensory evidence. This will induce action if the inferred state of the world carries with it proprioceptive predictions that are a poor match for the current proprioceptive input at peripheral levels. In this way, the distinction made in Section 4.1 between active and perceptual inference as different modes to reducing prediction error is artificial: the brain is always engaged in perceptual inference, with action occurring when its best estimate of the state of the world happens to entail proprioceptive prediction errors.

Prior expectations, furnished by sensory sampling, development, and evolution, are central to this process. For example, the nervous system may have such strong prior expectations that blood glucose is in a certain range (this prior being shaped by natural selection) that it is very unlikely to ever represent blood glucose as being outside of this range. A highly precise Bayesian prior for blood glucose levels is manifest in predictive processing as a prediction for interoceptive sensory activity together with attenuation of the weighting of interoceptive prediction errors (such that prediction errors drive autonomic reflexes, such as insulin release, or bodily actions). In this way, highly precise predictions will tend not to be updated in light of prediction error at the sensory level. Action is thus

---

hypothesized (and initially non-veridical) state of the world represented by motor predictions doesn't drive inference, and action can be initiated. The distinction between these accounts is discussed in Brown et al. (2013). Empirically, sensory attenuation has been examined in part by comparing the perception of self-generated vs. externally generated actions, for instance in force-matching and self-tickle paradigms (Brown et al., 2013; e.g., Van Doorn, Hohwy, & Symmons, 2014).

comparable to perceptual illusions in that prior expectations produce a non-veridical representation of the world; the difference is that in the case of action, peripheral reflexes proceed to rectify the mismatch between predictions and sensory data by making the real world fit with that modelled internally. In other words, returning to our previous example, our hand begins to move towards the coffee cup because that state of the environment is calculated, based on prior and contextual information, as being *more probable* than the (hidden) veridical state. This highlights the internalist nature of the brain under predictive processing: being compelled to infer a non-veridical model of the world is what enables an inferential system to act and maintain itself within expected states.

#### **4.4 The action-perception balance is central to autism**

As described in the preceding sections, precision estimation is fundamental to regulating the balance between perception and action, with implications for both perceptual inference and behaviour. The hypothesized differences in precision modulation in ASD thus suggest not only perceptual biases (as reviewed in Section 2.2) but also unusual patterns of behaviour. ASD is defined on the basis of impeded development in social interaction and communicatory behaviours (e.g., reduced social orienting), together with repetitive and restricted non-social behaviours (e.g., stereotyped movements and a strong preference for routine; American Psychiatric Association, 2013; Lai et al., 2014). These core aspects of ASD relate fundamentally to how individuals *interact with* or *sample* the world. Active inference points to how such interaction isn't a merely a matter of acting upon (potentially distorted) perceptual representations, but rather requires regulation of the *balance* between perceptual updating and action; relating, for instance, to where and for how long we sample in the visual field to optimise internal models, and how we interact with others to elicit information about their mental states. Investigating the outcomes of aberrant precision modulation for how action and sensory sampling are regulated may thus be crucial to moving from a Bayesian account of perceptual symptoms in ASD to an account of the core behavioural features that define the condition.

Reduced sensory attenuation (relating to chronically high precision expectations or impaired volatility processing, for example) is immediately suggestive of prolonged sampling behaviour at the expense of active investigation of the environment. Atypical *visual sampling behaviour* consistent with this hypothesis is a distinct characteristic of children with ASD, including in the years before diagnosis is typically made. This can take the form of unusual fixation on particular stimuli or object types (e.g., bright moving lights, carpet patterns, and small round objects), longer latencies in disengaging attention from one visual stimulus in order to orient to another, and unusual visual exploration of toys or toy-like objects (Bryson et al., 2007; Elsabbagh et al., 2009; Ozonoff et al., 2008; Zwaigenbaum et al., 2005). Similarly, visual exploration of object arrays shows a similar persistent pattern of gaze behavior, such that fewer objects are examined, yet with those examined being explored in more detail (Sasson, Elison, Turner-Brown, Dichter, & Bodfish, 2011; Sasson, Turner-Brown, Holtzclaw, Lam, & Bodfish, 2008). Features of this type are included in contemporary diagnostic criteria as “unusual interest in sensory aspects of the environment” (American Psychiatric Association, 2013), and an active area of investigation is whether the tendency towards a more persistent attentional style in ASD, occurring early in development, underlies the subsequent emergence of social abnormalities (Keehn, Muller, & Townsend, 2013).<sup>9</sup>

---

<sup>9</sup> Other findings are less clearly consistent with a bias against action initiation in autism, yet still implicate differences in the balance between perceptual updating and action as a central feature of autism early in development. A well-replicated psychophysical feature of autism is superior *visual search*, for which the location of a known target element embedded in an array of distractor elements is performed more quickly than in children without autism (Jarrold, Gilchrist, & Bender, 2005; Plaisted, O'Riordan, & Baron-Cohen, 1998). This finding has been replicated in adults (Kemner, van Ewijk, van Engeland, & Hooge, 2008; O'Riordan, 2004) and similarly extended to pre-verbal toddlers using gaze tracking (Kaldy, Kraper, Carter, & Blaser, 2011). The centrality of visual search behaviour to autism has been recently demonstrated in a prospective longitudinal study that tracked the younger siblings of individuals with autism, a cohort that are at increased risk of developing autism themselves. Spontaneous visual sampling behaviour was assessed implicitly with the use of gaze tracking when faced with a two-dimensional array of characters. Children whose gaze was drawn more readily to novel elements in these arrays at 9 months of age exhibited greater autistic symptoms at 15 months and 2 years of age (assessed using standard diagnostic inventories). A similar prospective study in infants found that heightened risk of

Similarly, there is evidence that *movement initiation* is impeded in ASD in the context of voluntary arm movements. Several studies report slower reaction times for young autistic adults in the initiation of cued pointing and pressing movements (Glazebrook, Elliott, & Lyons, 2006; Glazebrook, Elliott, & Szatmari, 2008; Nazarali, Glazebrook, & Elliott, 2009; for a conflicting finding, see Stoit, van Schie, Slaats-Willemse, & Buitelaar, 2013) and cued finger tapping (Turner, Frost, Linsenbardt, McIlroy, & Muller, 2006). Children with ASD similarly show longer and more variable preparation time before beginning movement of a stylus towards an illuminated target on a touch screen (Dowd, McGinley, Taffe, & Rinehart, 2012; Rinehart, Bellgrove, et al., 2006) as well as atypicalities in movement preparation in a button-pressing task (Rinehart, Bradshaw, Brereton, & Tonge, 2001; Rinehart, Tonge, et al., 2006). Longer overall durations of grasping and pointing movements (when excluding time before movement initiation) is similarly a well-replicated finding in adults and school-age children (Glazebrook et al., 2006; Glazebrook et al., 2008; Mari, Castiello, Marks, Marraffa, & Prior, 2003; Rinehart, Tonge, et al., 2006; Stoit et al., 2013). Palmer, Paton, Kirkovski, et al. (2015) examined the early constituent sub-components of movement in adults with ASD while they performed simple reach-to-grasp actions; early stages of movement were slow in this group, with delayed time to peak velocity of the initial sub-component of movement and delayed onset of later sub-components. When the process of movement initiation is considered in terms of active inference, these findings together suggest that the sensory attenuation required to facilitate a non-veridical hypothesis of arm position (and thus initiate movement) is impaired in ASD, reflecting an overweighting of sensory prediction error or atypical volatility expectations.

---

developing autism was associated with less time sampling from a specific part of two-dimensional static images (e.g., face, car) before shifting gaze to a different part of the image (Wass et al., 2015). Kemner, Verbaten, Cuperus, Camfferman, and van Engeland (1998) similarly find a greater number of saccades when viewing static two-dimensional images in school-age children with autism, comparing to both nonclinical children and children with ADHD. Whether effects of this nature can be modelled together with instances of persistent visual attention in a Bayesian framework is an interesting challenge to the theory.

In addition to a strong preference for routine, *repetitive or stereotyped behaviours* that occur in ASD include self-focussed motor behaviours (e.g., rocking or hand flapping) and repetitive use of objects (e.g., continually twirling a piece of string in front of the eyes). Prospective and retrospective studies indicate that atypical behaviours of this type can be apparent at 10-12 months of age in children that go on to receive a diagnosis of ASD (e.g., Bryson et al., 2007; Ozonoff et al., 2008; Werner, Dawson, Munson, & Osterling, 2005). These behaviours are similarly suggestive of a circumscribed sampling of the environment; as with tendencies towards prolonged visual sampling, repetitive movements are a way of continually sampling sensory information regarding the same external cause of input such that hypotheses are made more precise (in the latter case, hypotheses regarding bodily movement rather than visual objects). Similarly, by acting repetitively, the sensory input (whether proprioceptive, visual, or otherwise) is made more predictable and its external causes less volatile, such that uncertainty in the external causes of sensory input is reduced. In this way, repetitive sampling (whether repetitive motor behaviour or persistent visual attention) can be seen as a strategy for reducing uncertainty in the brain's representation of the world. As outlined in Section 4.2, the optimal way of sampling the environment to reduce uncertainty in perceptual representations is a fundamental challenge to perceptual inference that relies upon expectations for precision; in this respect, autistic symptoms may reflect a subtly different solution to this challenge, where maximal reduction in uncertainty in the brain's representation of the environment is linked to repeated sampling or movement rather than more variable or exploratory behaviours. Strategic differences in sampling behaviour can be cast in terms of volatility expectations; in a world that is less volatile, repetitive sampling may be a more optimal strategy to reducing uncertainty, while, in contrast, a stronger expectation that the causes of sensory input will change (including self-generated causes of sensory input) may facilitate less repetitive and exploratory movements.

The *socially* pertinent behaviours that define ASD in childhood and adulthood can similarly be considered in terms of how the sensory world is interactively sampled for the sake of perceptual inference. Early signs of ASD include reduced orienting

to social stimuli (e.g., an individual's response to their name being called, and the extent to which they attend preferentially to people rather than objects) and reduced joint attention (including both initiating shared attention with another by directing their attention using eye contact, gestures and speech, and following other's locus of attention based on their eye gaze and bodily gestures) (Bryson et al., 2007; Dawson et al., 2004; Zwaigenbaum et al., 2009). Atypical responses to direct gaze, including a reduced tendency to fixate on others' eyes (i.e., establish eye contact), has been long associated with ASD, with the current balance of evidence inconsistent but suggesting reduced direction of gaze towards others' eyes in more demanding and interactive contexts (i.e., potentially more volatile environments; Senju & Johnson, 2009). Similarly, autistic adults with intact ability to explicitly represent others' mental states (i.e., explicit theory of mind) show reduced eye gaze towards stimuli made implicitly salient by others' mental states (Senju, Southgate, White, & Frith, 2009). These differences in orienting sensory sampling towards social stimuli, and on the basis of others' mental states, suggest that socially related features of the environment, which readily capture attention, sampling and interactive behaviours in typically developing children and non-clinical adults, are afforded less precision in the brain of individuals with ASD and thus do not drive active inference to the same extent. This is suggestive of differences in the implicit counterfactual modelling of how the sensory consequences of actions that we perform are dependent on others' mental states, such that action is driven less by the imperative to reduce uncertainty in social aspects of the environment (Palmer, Seth, et al., 2015).

This section has argued that sensory-sampling behaviours, motor initiation, repetitive motor behaviours, and social interaction in ASD can be understood in terms of the balance between action and perceptual updating during inference. In predictive processing, the regulation of action rests upon the context-sensitive adjustment of precision weightings, itself depending on expectations about volatility and the deep hierarchical modelling of interacting worldly causes. Specifically, extended sampling behaviour at the expense of action, evidence for which exists in visual behaviour, slower movement initiation and repetitive movements, is suggestive of an expectation for less change in the environment.

This implicates mechanisms involved in hierarchical volatility inference and the recruitment of higher levels in the causal hierarchy to modulate predictions at lower levels.

## 5. Conclusion

This paper began by reviewing how simple and systematic differences in the processing of sensory information in the brain can be characterised within a Bayesian framework to encapsulate important perceptual characteristics of ASD. In the last several years, this approach has furnished initial hypotheses about ASD that can be considered across cognitive, computational and neural levels of description. We then explored theoretical challenges to optimal precision modulation in the brain, including the need for *volatility processing* and *hierarchical interactions* to maintain perceptual inference in the changing environments that the brain operates within. The promise of Bayesian theories of ASD is partly in how these more nuanced inferential mechanisms may allow for a rich account of context-dependency and heterogeneity in the expression of autistic characteristics. Moreover, the further expansion of the basic Bayesian model of perception, to include the interaction between perception and action (i.e., active inference), has promise for linking differences in the processing of sensory information in the brain to a broad set of autistic symptoms. Importantly, the hypothesised differences in the control of precision modulation in the autistic brain are directly suggestive not only of perceptual biases, but also unusual patterns of social and non-social behaviour.

## Linking text between chapters 1 and 2

In *Chapter 1*, predictive processing was introduced as a theory of how the brain builds a representation of its environment based on ambiguous sensory information. In the following chapters (*Chapters 2 and 3*) the processing of sensory information is examined in the context of *body perception*. We employ the *rubber hand illusion*, a paradigm that has been used widely in neuroscience research over the past decade to investigate how the brain represents the body as distinct from the external environment, the phenomenological sense of ownership we feel for our own body, and the underlying role of multisensory integration in these processes (Ehrsson, 2012).

To induce the illusion, the participant is typically seated in front of a prosthetic arm, with his or her own arm hidden from view. Ongoing touch, such as repeated stroking or brushing, is then applied synchronously to both the participant's hidden hand and the corresponding point of the prosthetic hand. *Crossmodal integration* of the visual and tactile information typically occurs, such that participants report the sensation of touch as spatially located where they see the prosthetic hand being touched – as if the tactile sensation is caused by the touch they see applied to the prosthetic hand. This illusion appears to exploit the fact that, in natural environments, touch to our body is commonly registered both visually and via tactile afferents, with temporal synchrony and anatomical congruence between the information that these two senses provide.

Most strikingly, this pattern of visual-tactile stimulation induces a strong sense that the prosthetic hand is a part of one's own body. This aspect of the illusion is most commonly assessed using self-report measures; however, several objective measures also support the notion that the illusion is associated with changes in the brain's representation of the body. Firstly, the illusion induces a shift in the sense of arm position towards the prosthetic arm, reflected in both explicit proprioceptive estimates and the trajectories of subsequent reaching movements (Botvinick & Cohen, 1998; Heed et al., 2011; Newport, Pearce, & Preston, 2010;

Tsakiris & Haggard, 2005; Zopf, Truong, Finkbeiner, Friedman, & Williams, 2011). Secondly, induction of the illusion is associated with changes in autonomic activity, including a drop in limb temperature (Hohwy & Paton, 2010; Kammers, Rose, & Haggard, 2011; Moseley et al., 2008; Thakkar, Nichols, McIntosh, & Park, 2011) and enhanced skin conductance responses when the prosthetic arm is threatened with physical damage (Armel & Ramachandran, 2003; Ehrsson, Wiech, Weiskopf, Dolan, & Passingham, 2007; Petkova & Ehrsson, 2009). Thirdly, *crossmodal congruency* effects indicate that visual stimuli presented nearby the prosthetic arm interfere with the processing of spatially-incongruent tactile stimulation of the real arm, reflected in measures of reaction time in a tactile discrimination task and suggesting that visual space around the prosthetic arm is in some regard ‘mapped’ to the body following the illusion (Pavani, Spence, & Driver, 2000; Zopf, Savage, & Williams, 2010).

Importantly, these phenomena tend to exist specifically when the seen and felt touch are applied synchronously, and are reduced for asynchronous stimulation. Studies that examine the constraints of the illusion have also shown that a degree of correspondence between the real and prosthetic arms in their position and orientation is necessary for the illusion to be induced, suggesting that congruence between *proprioceptive* and visual information also facilitates body representation (Costantini & Haggard, 2007; Hohwy & Paton, 2010; Lloyd, 2007). Neuroimaging studies of the illusion tend to implicate cortical regions that are known to integrate visual, tactile and proprioceptive signals, including the ventral premotor area and intraparietal regions (Brozzoli, Gentile, & Ehrsson, 2012; Ehrsson, Holmes, & Passingham, 2005; Ehrsson, Spence, & Passingham, 2004; Ehrsson et al., 2007). These findings thus support a model of body representation as resting in significant part on correlations between visual, tactile and proprioceptive information (discussed in Botvinick, 2004; Botvinick & Cohen, 1998; Ehrsson et al., 2004; Makin, Holmes, & Ehrsson, 2008).

In addition to multisensory correlations, what can be characterised as more *top-down* perceptual mechanisms play a role in body representation. Most notably, attempts to induce the rubber hand illusion using anatomically implausible

prosthetic arms or other non-bodylike objects, such as sticks and boxes, have shown absent or reduced effects (Hohwy & Paton, 2010; Tsakiris, Carpenter, James, & Fotopoulou, 2010; Tsakiris & Haggard, 2005). This suggests that implicit expectations regarding the visual appearance of the body contribute to body representation.

Influential theories of autistic perception propose differences in the interplay between higher- and lower-level perceptual mechanisms, with a reduced role for integrative processing of sensory information (Happé & Frith, 2006). There is also growing evidence for differences in the crossmodal integration of sensory information in ASD (e.g., reduced integration of visual and auditory information during speech perception; Mongillo et al., 2008). Effects of this nature are largely yet to be explored in the context of body perception and movement. The experience of the illusion may therefore be revealing about differences in the processing of sensory information in ASD, including that relevant to recent Bayesian hypotheses regarding autistic perception. For instance, the illusion entails conflict between proprioceptive signals regarding arm position and expectations for arm position influenced by the illusion. In this respect, proprioceptive and kinematic effects of the illusion may be revealing about the differences hypothesised for ASD in the weighting of incoming sensory information against prior or contextual expectations.

The RHI had been studied only very recently in ASD prior to the work reported in the present thesis. Paton et al. (2012) examined the illusion in a sample of high-functioning adults diagnosed with either autistic disorder or Asperger's disorder. Both clinical participants and non-clinical controls reported the typical subjective experience of the illusion. However, the ASD group were less sensitive to the degree of spatial discrepancy between the real and prosthetic arm in their experience of the illusion, and also exhibited subtle differences in the acceleration of reach-to-grasp movements performed with the stimulated hand following the illusion. The effects of the RHI have also been examined in children with ASD compared to typically developing children (ages 8–17; Cascio et al., 2012). The clinical group reported experiencing the typical subjective effects of the illusion,

but required a *longer duration* of visual-tactile stimulation before an effect on proprioceptive estimates was observed.

Together these initial findings suggest intact integration of seen and felt touch in ASD, but with reduced interaction of this stimulation with proprioceptive representations. These results cohere partly with evidence that individuals with ASD exhibit a greater reliance than controls on proprioceptive input rather than visual information when guiding reaching movements (Haswell, Izawa, Dowell, Mostofsky, & Shadmehr, 2009; Masterton & Biederman, 1983) (see also Glazebrook, Gonzalez, Hansen, & Elliott, 2009) and maintaining postural stability (Minshew, Sung, Jones, & Furman, 2004).

In *Chapter 2*, susceptibility to the rubber-hand illusion is examined in a sample of nonclinical adults. Proprioceptive and motoric effects of the illusion are found to differ in this sample with respect to variation in autism-like traits. These findings are then partially replicated and extended in *Chapter 3*, which reports data from adults with formal ASD diagnoses in addition to nonclinical individuals that vary in their autism-like traits. *Chapter 3* also characterises the motoric effects of the illusion in greater detail; we find that it is not the *initiation* of movement that is influenced by the illusion, but rather *later* stages of movement, consistent with a difference in ASD in the integration between sensory signals received during movement and expectations for arm position influenced by the illusion. These findings are discussed in the context of predictive processing theories of neurocognitive function.

## Declaration for Thesis Chapter 2

### Declaration by candidate

In the case of *Chapter 2*, the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
<ul style="list-style-type: none"> <li>Contributed to the experiment design.</li> <li>Performed data collection and data analysis.</li> <li>Contributed to the interpretation of results.</li> <li>Wrote the paper.</li> </ul>	50%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Bryan Paton	Technical input to data collection and analysis. Contributed to the experiment design, interpretation of results, and writing.	15%
Jakob Hohwy	Contributed to the experiment design, interpretation of results, and writing.	N/A
Peter Enticott	Contributed to the experiment design, interpretation of results, and writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date 16/02/2016
--------------------------	---	--------------------

Main Supervisor's Signature		Date 16/02/2016
-----------------------------------	---	--------------------

## Chapter 2

# Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum

Colin J. Palmer<sup>1</sup>, Bryan Paton<sup>1,2</sup>, Jakob Hohwy<sup>1</sup> and Peter G. Enticott<sup>3</sup>

<sup>1</sup>*Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup>*School of Psychology & Psychiatry, Monash University, Clayton, VIC 3800, Australia*

<sup>3</sup>*Monash Alfred Psychiatry Research Centre, The Alfred and Central Clinical School,  
Monash University, Melbourne, VIC 3004, Australia*

Published in *Neuropsychologia*, 2013

## Abstract

Recent research has begun to investigate sensory processing in relation to nonclinical variation in traits associated with autism spectrum disorder (ASD). We propose that existing accounts of autistic perception can be augmented by considering a role for individual differences in top-down expectations for the precision of sensory input, related to the processing of state-dependent levels of uncertainty. We therefore examined autistic traits in relation to the rubber-hand illusion: an experimental paradigm that typically elicits crossmodal integration of visual, tactile, and proprioceptive information in an unusual illusory context. Individuals with higher autistic traits showed reduced effects of the rubber-hand illusion on perceived arm position and reach-to-grasp movements, compared to individuals with lower autistic traits. These differences occurred despite both groups reporting the typical subjective experience of the illusion concerning visuotactile integration and ownership for the rubber hand. Together these results suggest that the integration of proprioceptive information with cues for arm position derived from the illusory context differs between individuals partly in relation to traits associated with ASD. We suggest that the observed differences in sensory integration can be best explained in terms of differing expectations regarding the precision of sensory estimates in contexts that suggest uncertainty.

## 1. Introduction

ASD frequently involves atypical sensory processing in both childhood and adulthood (reviewed in Iarocci & McDonald, 2006; Marco et al., 2011; Simmons et al., 2009). The upcoming fifth edition of the Diagnostic and Statistical Manual of Mental Disorders will for the first time include sensory dysfunction as a diagnostic criterion for ASD (i.e., “hyper- or hyporeactivity to sensory input or unusual interest in sensory aspects of the environment,” American Psychiatric Association, 2013), calling attention to the need to advance our understanding in this area. To understand the nature of ASD and to throw light on individual differences in perception more generally, it is also important to explore the extent to which the relevant underlying sensory mechanisms vary in the general population.

This broader focus of research originates from evidence that autistic traits vary meaningfully amongst nonclinical individuals, with those meeting a clinical diagnosis of ASD situated at the extreme end of a spectrum that encompasses the population at large (reviewed in Happé et al., 2006; Mandy & Skuse, 2008). The distribution of scores typically found for measures of autistic traits in large general population samples tends to be compatible with this hypothesis (e.g., Constantino & Todd, 2003; Posserud, Lundervold, & Gillberg, 2006), and correlations between autistic traits and sensory task performance in non-clinical samples are consistent with sensory differences seen in clinically-diagnosed ASD (e.g., Donohue, Darling, & Mitroff, 2012; Walter et al., 2009). A similar technique used to investigate phenomena related to ASD is the group comparison of nonclinical individuals scoring high on trait measures of ASD to those scoring lower. This approach has also revealed sensory differences (Grinter, Maybery, et al., 2009; Grinter, Van Beek, Maybery, & Badcock, 2009) and neurophysiological response characteristics (Puzzo, Cooper, Vetter, & Russo, 2010) associated with autistic traits consistent with that seen in clinically-diagnosed ASD, and this method is employed in the present study.

Contemporary theories of perception in ASD propose fundamental differences in the processing of sensory information to account for a complex pattern of

strengths and weaknesses observed across different perceptual-cognitive tasks and contexts (e.g., Brock, Brown, Boucher, & Rippon, 2002; Frith, 1989; Happé & Frith, 2006; Mottron et al., 2006; Plaisted et al., 1998). A theme underlying parts of this discussion, in particular the weak central coherence theory (Frith, 1989; Happé & Frith, 2006), is the neurocognitive distinction between the contribution of bottom-up sensory processing to perception (relating most directly to sensory input) and the top-down modulation of input based on endogenous factors such as prior knowledge and attention (Frith & Dolan, 1997; Gilbert & Sigman, 2007; Kveraga, Ghuman, & Bar, 2007). More recent Bayesian accounts develop this point in relation to ASD explicitly: for example, Pellicano and Burr (2012) suggest that prior expectations regarding the state of the world may have diminished influence on perception in ASD, increasing reliance on bottom-up signals (for discussion and related proposals, see Brock, 2012; Friston, Lawson, & Frith, 2013; Hohwy, 2013; Mitchell & Ropar, 2004; Paton et al., 2012; van Boxtel & Lu, 2013b).

An important challenge for these accounts is the uneven landscape of enhanced and compromised perceptual performance in ASD, which does not cohere clearly with a general bias in top-down processes. For example, for visual illusions, some, but not all, studies have suggested less susceptibility (that is, increased veridical perception) in ASD (Bölte et al., 2007; Happé, 1996; Hoy et al., 2004; Ropar & Mitchell, 1999, 2001; Walter et al., 2009). Similarly, whereas a general impairment in top-down modulation would seem to predict diminished multisensory integration in ASD, studies do not unequivocally support this, even though there are a number of intriguing underlying differences (Cascio et al., 2012; Kwakye, Foss-Feig, Cascio, Stone, & Wallace, 2011; reviewed in Marco et al., 2011; Paton et al., 2012).

We reasoned that uneven performance could relate to differences in the way context determines the *expected levels of sensory precision*, which is an aspect of top-down modulation that has only recently been described (Feldman & Friston, 2010) and linked to ASD in the context of predictive processing models of perception (Friston et al., 2013; Paton et al., 2012). Conceptually, expectations regarding the precision of sensory input are of importance to the relative

weighting of bottom-up and top-down perceptual processes in response to state-dependent (i.e., changing) levels of uncertainty. This proposal therefore predicts that differences will become apparent in contexts and experimental set-ups where changing conditions suggest changing levels of uncertainty in the sensory signal. In particular, individuals with ASD, as well as nonclinical individuals with autistic traits, may be less sensitive than individuals with few autistic traits to contexts that suggest increased uncertainty. This would predict that in contexts that suggest low uncertainty (i.e., high precision of sensory input) there would be less difference between the groups, but that in contexts that suggest higher uncertainty (i.e., suggests low precision of sensory input) differences would begin to emerge. Sometimes these differences would give rise to enhanced performance of the ASD and autistic groups, namely when the expectation for high precision input leads to increased sensory sampling and less integration under prior expectations relevant to the context. Sometimes this would lead to compromised performance for these groups, namely when expectation for high precision leads to blindness to underlying patterns of hidden, influencing factors.

We therefore explore this proposal in relation to the rubber-hand illusion (Botvinick & Cohen, 1998), a well-studied experimental paradigm involving multisensory interactions in relation to the neural representation of body location. Here, repetitive tactile stimulation is applied synchronously to the participant's hand (hidden from view) and a fake rubber hand (that lies in view). This pattern of sensory input typically induces the illusory sensation that touch is felt on the surface of the rubber hand, as well as a heightened sense of ownership for the rubber hand (see Ehrsson, 2012, for review). The integration between visual and tactile sensory inputs is also associated with a measurable drift in perceived hand location towards the rubber hand (Botvinick & Cohen, 1998; Tsakiris & Haggard, 2005) and subtle changes in subsequent reaching movements performed with the stimulated hand (e.g., Kammers, Kootker, Hogendoorn, & Dijkerman, 2010). Importantly, these phenomena tend to exist specifically when the seen and felt touch are applied synchronously, and are reduced for asynchronous stimulation. This paradigm therefore involves both sensory integration under different global

causal models and, also, a highly unusual, uncertainty-inducing context of experiencing touch on a rubber hand.

As described above, the rubber-hand illusion should be expected to trigger differences in expected precisions of visual, tactile and proprioceptive sensory input. Our previous study of this illusion (Paton et al., 2012) compared a clinical ASD group with healthy controls and found differences in proprioception and motor parameters on a reach task. Following the results of this study, we expect that participants will experience the typical subjective effects of the illusion (e.g., that touch is mislocated to the rubber hand) regardless of their level of autistic traits, and thus rate the strength of these effects, as assessed via questionnaire, stronger during synchronous than asynchronous stimulation. We further predict that individuals with autistic traits will show less sensitivity to the presence of the illusion in their perceived arm position than individuals low on autistic traits (i.e., less of a difference in proprioceptive drift between synchronous and asynchronous stimulation conditions). This hypothesis is based on the notion of lower sensitivity to state-dependent uncertainty in individuals with autistic traits, and coheres with the previous finding of more accurate proprioception in individuals with ASD compared to controls (Paton et al., 2012).

In addition, it is predicted that reaching movements performed subsequent to the illusion will reflect the uncertainty suggested by the unusual illusory content. This latter hypothesis is based on the idea that expectations regarding the precision of sensory (proprioceptive) input occurring as movement unfolds affect how smoothly movement is performed. In short, if proprioceptive imprecision is expected, movement should be uncertain, exploratory, and tentative (cf. Friston et al., 2010). Specifically, we expect that individuals with low autistic traits will exhibit less smooth movement after experiencing the illusion than individuals with high autistic traits. Higher order temporal derivatives of position (e.g., jerk) are of interest to this hypothesis due to their relationship with movement smoothness. Our previous study, which found differences between clinical ASD and control participants in the acceleration of reaching movements performed following the illusion, was unable to assess comprehensively differences in movement (such as

smoothness) due to limits of the tracking technique used. The current study therefore extends previous findings to a nonclinical sample of individuals with and without autistic traits and asks, in particular, whether the differences in motor parameters could pertain to differences in expected precisions.

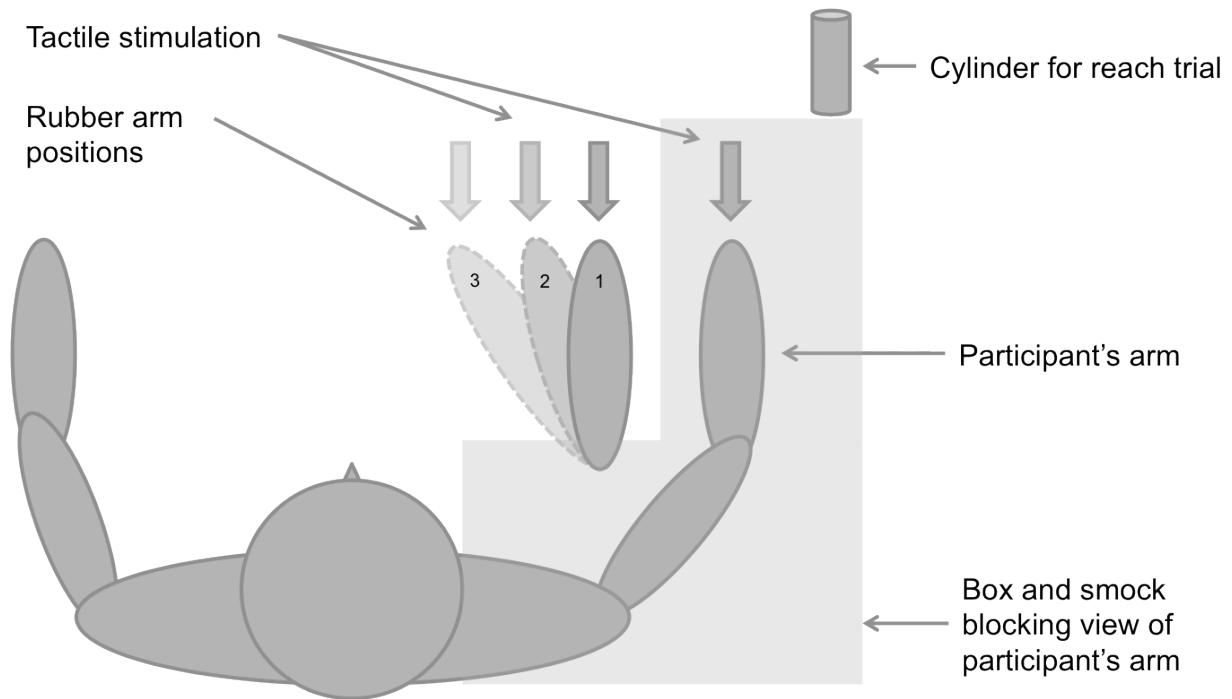
## 2. Method

### 2.1 Participants

Twenty-four right-handed individuals ( $M = 28.96$ ,  $SD = 11.16$  years; 13 female) completed the experiment. Volunteers were recruited via advertisements distributed to the general Monash University population. Participants were separated into two groups based on a median-split of their scores on the Autism-Spectrum Quotient (AQ, described below; whole sample:  $M = 116.33$ ,  $SD = 14.47$ ; low AQ group:  $M = 104.58$ ,  $SD = 8.95$ ; high AQ group:  $M = 128.08$ ,  $SD = 7.5$ ). Note that, while the present study used Likert scoring for the AQ, values for the AQ using binary scoring are as follows: whole sample,  $M = 21.33$ ,  $SD = 6.83$ ; low AQ group,  $M = 16$ ,  $SD = 3.84$ ; high AQ group,  $M = 26.67$ ,  $SD = 4.56$ . Each group contained 12 participants (low AQ group:  $M = 32.42$ ,  $SD = 13.98$  years, 7 female; high AQ group:  $M = 25.50$ ,  $SD = 6.26$  years, 6 female). The study was approved by the Monash University Human Research Ethics Committee. All participants provided written informed consent.

### 2.2 Materials and Procedure

Participants were seated in accordance with the experimental set-up illustrated in *Figure 1*.



**Figure 1.** Experimental set-up. The rubber arm and participant's right arm were placed in separate compartments. A prosthetic right limb was used with a high degree of visual similarity to a human limb regarding physical dimensions, skin detail, and compression to touch. A semi-silvered mirror lid enabled the experimenter to control the participant's vision into either compartment via adjustment of the lighting inside. Participants were able to see the rubber arm only during the stimulation phase of each trial. The participant's own arm was occluded from view throughout the experiment. The cylindrical reach target was only visible during the reaching phase of each trial.

### 2.2.1 Independent variables

The position of the rubber arm was varied between three positions across trials. Synchronous and asynchronous tactile stimulation was delivered independently for each position of the rubber arm. Trials were conducted in two blocks, each comprised of a single trial for each of the six conditions. Participants therefore completed two trials for each of the six conditions, and dependent measures were averaged across these two trials. To control for order effects, trial order was

randomised for each block across participants. The duration of the experiment was 90–120 min.

**2.2.1.1 Stimulation type.** An experimenter seated opposite to the participant manually applied repetitive tactile stimulation to anatomically corresponding locations of the participant's right hand and the rubber arm. Stimulation was applied at approximately 1–2 Hz for 3 min in each trial with a pair of small paintbrushes (2–2.5 x 0.5 cm brush area). Trials involved either synchronous or asynchronous stimulation for the entire period. Asynchronous tactile stimulation is typically used as a control condition in rubber-hand illusion studies, as temporal synchronicity between the seen and felt touch is associated with significantly stronger perceptual effects (Botvinick & Cohen, 1998; estimated as best within approximately 300 ms, Shimada, Fukuda, & Hiraki, 2009). Stimulation during the asynchronous condition was both temporally and spatially asynchronous. Participants were directed to attend to the rubber hand during the stimulation period.

**2.2.1.2 Rubber arm position.** The participant's arm rested in the same position for every condition. The rubber arm was varied between three positions such that the horizontal distance separating the middle finger of each hand was 20 cm, 25 cm, or 30 cm (numbered 1–3 in in *Figure 1*, and referred to as positions 1–3 henceforth). The orientation of the rubber arm changed between each position such that the end proximal to the participant always entered the box in line with the participant's right shoulder. This was intended to maintain anatomical plausibility for ownership of the rubber arm across conditions, which is a known constraint on illusion induction (Ehrsson et al., 2004; Pavani et al., 2000; Tsakiris & Haggard, 2005). The orientation of tactile stimulation on the rubber hand was adjusted across positions to maintain congruency in the direction of stimulation applied to the real and rubber hands in a hand-centred reference frame (see Costantini & Haggard, 2007, for an investigation of orientation mismatch in hand-centred versus external space reference frames).

Anatomical congruence between the placement of the rubber arm and the position of the real arm has been shown to influence the strength of the rubber-hand illusion (Costantini & Haggard, 2007; Ehrsson et al., 2004; Ide, 2013; Lloyd, 2007; Pavani et al., 2000; Tsakiris & Haggard, 2005; see also White & Aimola Davies, 2011). These findings may relate to expected precisions, in as much as different positions of the (real) arm have been shown to have different proprioceptive precisions (van Beers et al., 1998). The effects of manipulating rubber arm position is of interest in relation to examining the influence of top-down processes on the illusory experience; for example, top-down processes comparing expectations regarding body position to that of the rubber arm (Tsakiris & Haggard, 2005). Increasing the distance of the rubber arm from the real arm was expected to decrease the self-rated strength of the illusion, as has been found previously in a nonclinical sample (Lloyd, 2007). We further hypothesised that individuals with stronger autistic traits may be less sensitive to changes in the anatomical congruence between the real and rubber arms than individuals with lower autistic traits, due to a lesser influence of top-down processes on perception. We therefore expected the latter group to be more likely to show differences in self-rated illusion strength and proprioceptive drift between the rubber arm positions during synchronous stimulation.

### 2.2.2 Dependent measures

Several dependent measures were collected in each trial to capture perceptual and sensorimotor effects of the rubber-hand illusion. Estimates of arm location were recorded directly before and after each stimulation period. A reach-to-grasp movement was conducted following the post-stimulation estimate of arm location. At the end of each trial, participants completed a questionnaire related to their subjective experience of the illusion. A psychological inventory designed to assess autistic traits (the AQ) was completed during a break midway through the twelve trials of the rubber hand illusion.

**2.2.2.1 Illusion ratings.** Participants completed a short questionnaire to report on their experiences during tactile stimulation. This consisted of 11 items displayed in

*Table 1*, adapted from Botvinick and Cohen (1998; I1–I3, C2, C4, C5), Longo, Schuur, Kammers, Tsakiris, and Haggard (2008; C6, C7), Petkova and Ehrsson (2008; C3), and Hohwy and Paton (2010; C1, C8). Three items (I1–I3) were statements describing the content of the illusion typically reported in the literature. Eight items (C1–C8) were included to control for response biases, and described possible experiences that were not expected to differ consistently between synchronous and asynchronous stimulation. Each item was rated on a 20 cm horizontal visual analogue scale with left and right endpoints marked as *strongly disagree* and *strongly agree*, respectively. The centre of the scale was labelled *very unsure whether agree or disagree*. Participants could mark the scale anywhere along its length, and markings were scored to the nearest millimetre. Greater values indicated stronger agreement with the statement. Items were presented in a fixed order across trials and in pen-and-paper format. Participants were required to remove their right arm from the box when completing this measure to help disrupt the effects of the illusion between trials.

**Table 1.** Self-rated illusion questionnaire.

Item type	No.	Text
Illusion	I1	It seemed as if I was feeling the touch of the paintbrush in the location where I saw the rubber hand being touched.
	I2	It seemed as though the touch I felt was caused by the paintbrush I could see touching the rubber hand.
	I3	It felt as if the rubber hand was my hand.
Control	C1	It felt as if my (real) hand was getting cold.
	C2	It seemed as if I might have more than one right hand or arm.
	C3	It seemed as if I was in two different locations at the same time.
	C4	It felt as if my (real) hand was turning 'rubbery'.
	C5	The rubber hand began to resemble my own (real) hand, in terms of shape, skin tone, freckles or some other visual feature.
	C6	I found the touch of the paintbrush on my hand was pleasant.
	C7	I found myself liking the rubber hand.
	C8	I felt the room temperature change during the experiment.

**2.2.2.2 Proprioceptive drift.** Participants were asked to estimate the position of their visually occluded right hand directly before and after each period of stimulation. For this procedure the experimenter slid a plexiglass marker across a rail that ran horizontally with respect to the participant along the top of the box. The participant verbally indicated when a vertical line on the marker was judged as being directly above the centre knuckle of their right hand. The location of the marker was recorded to the nearest millimetre via a fixed ruler (only visible to the experimenter) that spanned the length of the rail. Participants were unable to see either the rubber arm or their own arm during this stage of each trial, and were asked to keep their arm still to limit proprioceptive feedback. A measure of *proprioceptive drift* was calculated for each trial by subtracting the participant's pre-stimulation estimate of hand location from their post-stimulation estimate. Positive values indicate that the estimate of hand location was closer to the rubber arm following stimulation. As with questionnaire ratings, proprioceptive drift is a common measure of illusion induction (e.g., Botvinick & Cohen, 1998; Tsakiris & Haggard, 2005). There is evidence distinguishing the neural substrates (Brozzoli et al., 2012; Ehrsson et al., 2004; Fiorio et al., 2011; Kammers, Verhagen, et al., 2009) and behavioural coincidence (Holmes, Snijders, & Spence, 2006; Rohde, Di Luca, & Ernst, 2011) of these measures, however, suggesting a distinction between the mechanisms underlying changes in perceived hand location and the subjective experience of ownership and tactile mislocation induced in the rubber-hand illusion set-up.

**2.2.2.3 Reach-to-grasp movement.** Following the post-stimulation estimate of hand location, participants were asked to reach out and grasp a cylinder with the hand involved in the stimulation period. The cylinder measured 4.5 cm diameter by 18 cm height and was located within the box 13 cm in front and 5 cm to the right of the participant's hand. Participants were able to see approximately the upper 2 cm of the cylinder during this phase of the experiment while both their arm and the rubber arm were occluded from view. Participants were instructed that they were not required to minimise their reaction time or maximize their speed of movement.

Hand trajectories were recorded using an electromagnetic tracker (Ascension Technology Corporation 3DGuidance trakSTAR with mid-range transmitter; 1.4 mm and 0.5 degrees static accuracy in an optimal environment). The six dimensions of translation and rotation were recorded via a magnetic sensor attached to the centre of the dorsal surface of the participant's right hand. These data were filtered with a 50 Hz notch filter to remove AC line noise and a third order zero-phase low-pass Butterworth filter with a cutoff at 20 Hz.

Participants were instructed to begin the movement when a light was switched on to allow vision of the target object. Position data were recorded continuously (60 Hz sample rate) for 5 seconds following this point. In an adaptation of the method used by Kammers, de Vignemont, Verhagen, and Dijkerman (2009) and Kammers, Verhagen, et al. (2009), movement onset was defined as when velocity first exceeded 20 mm/s continuously for 0.05 seconds. Movement offset was defined as when velocity first exceeded 20 mm/s for 0.05 seconds when proceeding retrograde through the time series. Twelve trials were discarded due to recording malfunction or on account of the participant failing to execute the movement as instructed.

The kinematic parameter of primary interest was the normalised integrated jerk of the reaching movement. Jerk is the derivative of acceleration with respect to time (i.e., the third derivative of position, mm/s<sup>3</sup>), and is commonly employed as a measure of movement smoothness (Hogan & Sternad, 2009). In the present study, integrated jerk was calculated as the area under the curve of the Euclidean jerk vector obtained from the three linear axes. Following previous research that has studied the jerk of voluntary movements in clinical populations, the integrated jerk for each trial was normalised for both movement extent and movement duration before undergoing analysis (Hogan & Sternad, 2009; Nobile et al., 2011; Romero, Van Gemmert, Adler, Bekkering, & Stelmach, 2003; Teulings, Contreras-Vidal, Stelmach, & Adler, 1997).

The present study also analysed several other kinematic parameters that have been examined in previous studies of reaching movement in the rubber-hand

illusion (Kammers, de Vignemont, et al., 2009; Paton et al., 2012; Zopf et al., 2011). Movement duration is the time between movement onset and offset as defined above. Mean velocity is the mean of Euclidean velocity across the time series. Peak velocity is the maximum instantaneous Euclidean velocity recorded across the time series. Relative time to peak velocity is the time between movement onset and when peak velocity is achieved, as a percentage of total movement duration. Integrated acceleration was calculated from the recorded trajectories as the area under the curve of the Euclidean acceleration vector. Peak horizontal displacement is the maximum of displacement in the horizontal dimension (with respect to the participant) in the direction towards the reach object and away from the rubber arm. Following Zopf et al. (2011), the angle of initial movement was calculated from the instantaneous velocity at the time point when 10% of the Euclidean displacement towards the end point was achieved. These latter two measures of hand displacement are of particular interest given that the rubber-hand illusion affects perceived hand position, which may be expected to influence the initial displacement of subsequent reaching movements towards a fixed target (Heed et al., 2011; Newport et al., 2010; Zopf et al., 2011). Hence, by analysing both displacement measures and other parameters (such as integrated jerk and movement duration) we hoped to distinguish to an extent between an effect of proprioceptive drift on subsequent reaching movements and other potential effects of the rubber-hand illusion on reaching movements.

**2.2.2.4 AQ.** The AQ is a self-administered and non-diagnostic 50-item questionnaire, designed to measure traits associated with ASD in adults (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001). The psychometric properties of this scale have received support for use in non-clinical participants with normal IQ (e.g., Baron-Cohen et al., 2001; Hurst, Mitchell, Kimbrel, Kwapił, & Nelson-Gray, 2007; Stewart & Austin, 2009). Each item consists of a statement; for example, “I prefer to do things the same way over and over again”. Participants rate their level of agreement with each statement on a 4-point Likert scale ('definitely agree', 'slightly agree', 'slightly disagree', 'definitely disagree') and responses are summed with 26 items reverse scored. The range of possible scores is 50–200, with higher scores indicating greater similarity to traits of ASD. This

approach to scoring differs from the binary system used by Baron-Cohen et al. (2001). Likert scoring is preferred in the current study to increase sensitivity to individual differences between nonclinical participants. This method of scoring has been used previously for the AQ (e.g., Stewart, Watson, Allcock, & Yaqoob, 2009), and there is evidence that Likert scoring is associated with improved psychometric properties compared to binary scoring for personality questionnaires (Muñiz, García-Cueto, & Lozano, 2005). In the present study, a pen-and-paper version of this scale was administered.

### **2.3 Statistical analyses**

The present study employed a mixed factorial design. The within-subjects factors were questionnaire item type (illusion items vs. control items), stimulation type (synchronous stimulation vs. asynchronous stimulation), and rubber arm position (position 1 vs. position 2 vs. position 3). The between-subjects factor was AQ group (low AQ group vs. high AQ group). Mixed between-within subjects ANOVAs were conducted for each dependent measure to assess for main and interaction effects across conditions. Normalised integrated jerk values were non-normally distributed in the present study (as has been found previously in the literature; Teulings et al., 1997). The nonparametric Wilcoxon signed-rank test was therefore used to examine for differences in jerk between the two stimulation conditions separately for each AQ group. Post-hoc tests were performed using Bonferroni correction to control for Type I error. Effect sizes are reported here using Cohen's *d*.

## **3. Results**

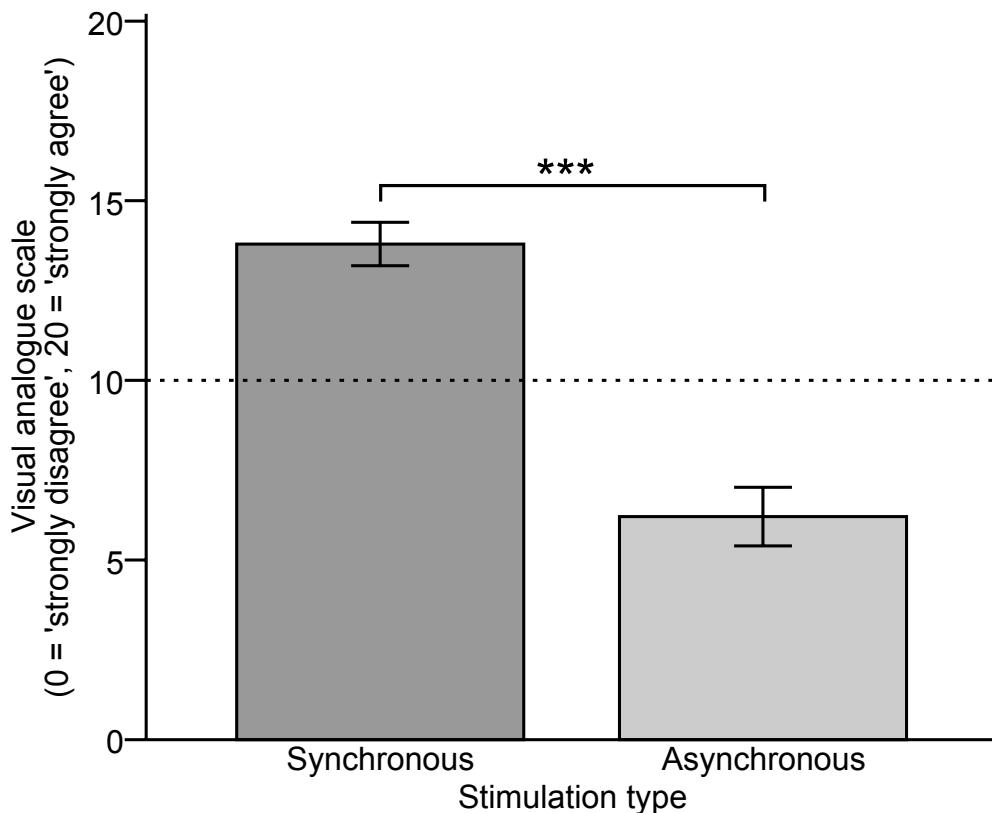
### **3.1 Illusion ratings**

A 2x2x2x3 mixed ANOVA was performed for illusion ratings with Group (Low AQ vs. High AQ) as a between-subjects factor and Item Type (Illusion vs. Control),

Stimulation Type (Synchronous vs. Asynchronous) and Rubber Arm Position (Position 1 vs. Position 2 vs. Position 3) as within-subjects factors.

A main effect was found for Item Type, indicating that illusion items ( $M = 10.00, SD = 2.81$ ) were rated higher than control items ( $M = 7.20, SD = 2.94$ ),  $F(1, 22) = 34.70, p < .001$ , Cohen's  $d = 0.97$ . A main effect was also found for Stimulation Type, indicating that item ratings were higher following synchronous stimulation ( $M = 10.80, SD = 2.48$ ) than following asynchronous stimulation ( $M = 6.40, SD = 3.27$ ),  $F(1, 22) = 79.81, p < .001$ , Cohen's  $d = 1.52$ . Importantly, an interaction effect was found between Item Type and Stimulation Type,  $F(1, 22) = 76.93, p < .001$ . Interaction effects were also found between Stimulation Type and Group,  $F(1, 22) = 4.53, p < .05$ , and between Item Type, Stimulation Type, and Group,  $F(1, 22) = 6.64, p < .05$ . There were no other significant main or interaction effects ( $p > .05$ ).

As expected, post-hoc tests indicated that synchronous stimulation ( $M = 13.79, SD = 2.96$ ) was associated with higher illusion item ratings than asynchronous stimulation ( $M = 6.21, SD = 4.00$ ),  $t(23) = 8.78, p < .001$ , Cohen's  $d = 2.15$  (*Figure 2*). Furthermore, synchronous stimulation was associated with significantly higher ratings on illusion items ( $M = 13.79, SD = 2.96$ ) than control items ( $M = 7.80, SD = 3.18$ ),  $t(23) = 8.09, p < .001$ , Cohen's  $d = 1.95$ . In contrast, there was no significant difference between ratings for illusion and control items for asynchronous stimulation,  $t(23) = -0.72, p = .48$ , Cohen's  $d = -0.11$ . Together these results indicate that synchronous stimulation induced the phenomenological features of the illusion typically reported in the literature more strongly than asynchronous stimulation. Further post-hoc analyses are presented in the Supplementary Material.



**Figure 2.** Ratings of illusion items across stimulation type conditions. Error bars indicate  $\pm 1$  standard error. (\*\*\*( $p < .001$ )).

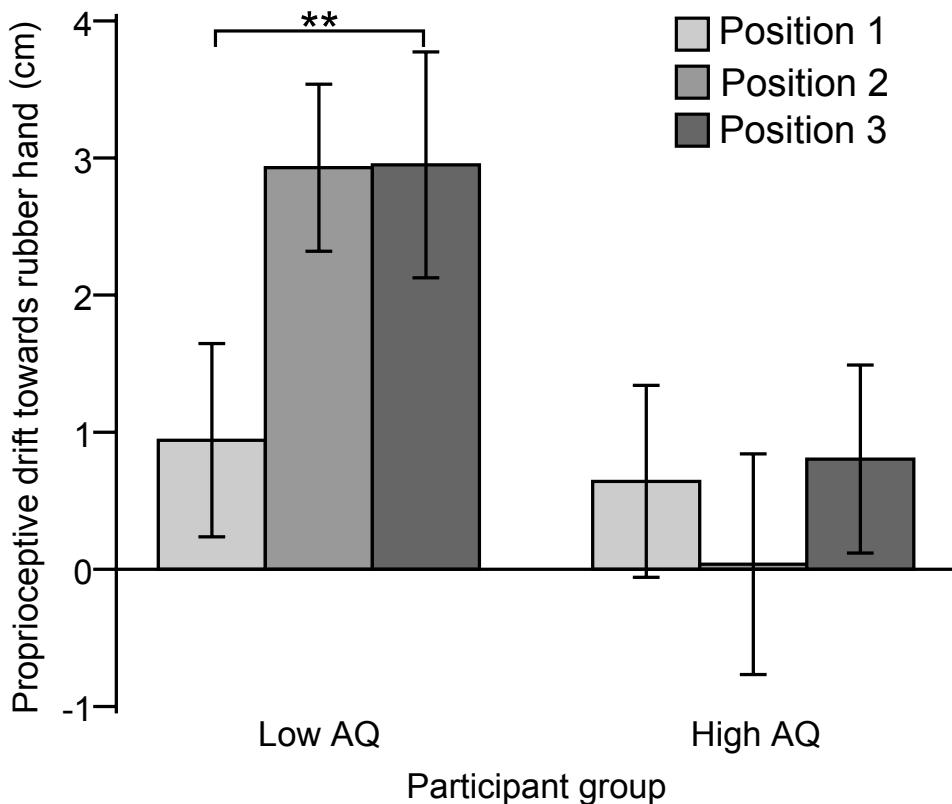
### 3.2 Proprioceptive drift

A 2x2x3 mixed ANOVA was performed for proprioceptive drift with Group (Low AQ vs. High AQ) as a between-subjects factor and Stimulation Type (Synchronous vs. Asynchronous) and Rubber Arm Position (Position 1 vs. Position 2 vs. Position 3) as within-subjects factors.

A significant main effect of Group was found for proprioceptive drift,  $F(1, 22) = 4.99, p < .05$ , Cohen's  $d = 0.91$ . The low AQ group displayed greater proprioceptive drift across conditions ( $M = 1.36, SD = 1.72$ ) compared to the high AQ group ( $M = -0.04, SD = 1.31$ ). A significant main effect of Stimulation Type was found for proprioceptive drift,  $F(1, 22) = 10.92, p < .01$ , Cohen's  $d = 0.73$ , indicating that synchronous stimulation ( $M = 1.38, SD = 2.19$ ) was associated with greater drift in

perceived arm position towards the rubber arm than asynchronous stimulation ( $M = -0.07, SD = 1.73$ ). There was no significant interaction between Stimulation Type and Group ( $p > .05$ ); however, a significant three-way interaction was found between Stimulation Type, Rubber Arm Position, and Group,  $F(2, 44) = 3.28, p < .05$ . No other main or interaction effects were found for this variable ( $p > .05$ ).

To clarify the three-way interaction effect, two-way repeated-measures ANOVAs were conducted separately for each AQ group, with Stimulation Type and Rubber Arm Position as factors. A significant Stimulation Type by Rubber Arm Position interaction effect was found for the low AQ group,  $F(2, 22) = 6.05, p < .01$ , but not the high AQ group,  $F(2, 22) = 0.19, p = .83$ . Further one-way ANOVAs for the low AQ group indicated a significant main effect of Rubber Arm Position for synchronous stimulation,  $F(2, 22) = 5.20, p < .05$ , but not asynchronous stimulation,  $F(2, 22) = 0.89, p = .43$ . Post-hoc tests for the low AQ group during synchronous stimulation indicated that significantly greater drift was observed for position 3 (30 cm separation between participant's arm and the rubber arm;  $M = 2.95, SD = 2.85$ ) compared to position 1 (20 cm separation;  $M = .94, SD = 2.44; p < .01$ , Cohen's  $d = 0.76$ ). To summarise, the degree of drift in arm position towards the rubber arm following synchronous stimulation was influenced by the distance of the rubber arm from the participant's arm, but only for the group of participants who scored lower on the AQ (*Figure 3*).



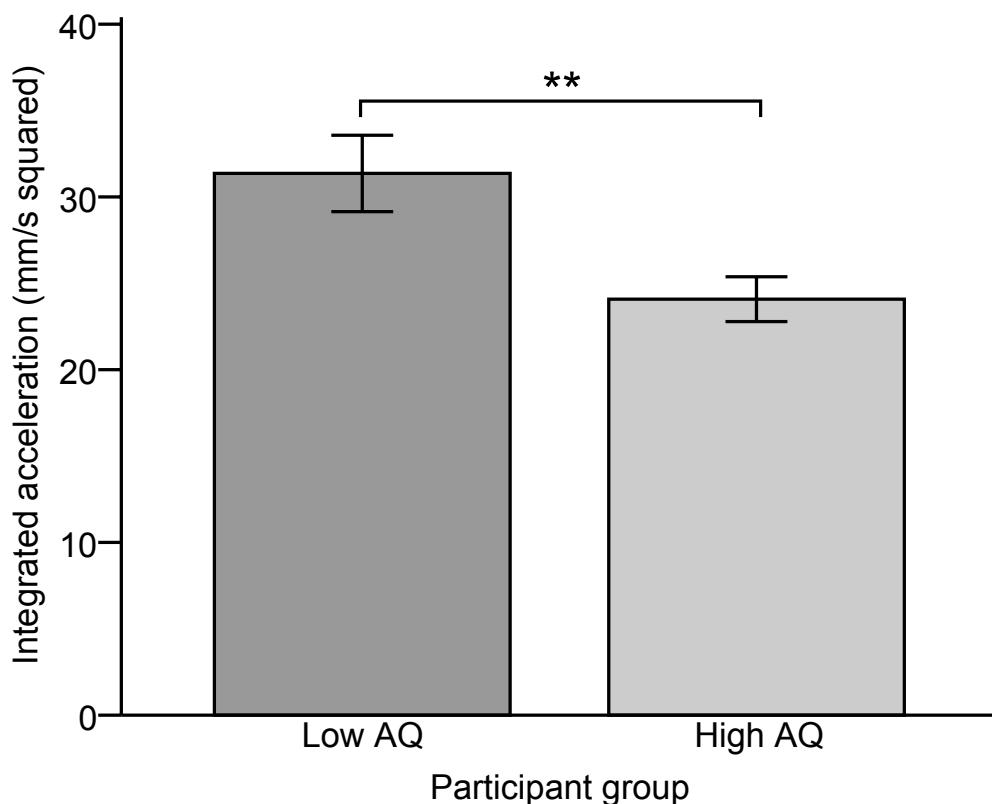
**Figure 3.** Proprioceptive drift towards the rubber hand across rubber arm position conditions following synchronous stimulation. Error bars indicate  $\pm 1$  standard error. (\*\* $p < .01$ ).

### 3.3 Integrated acceleration

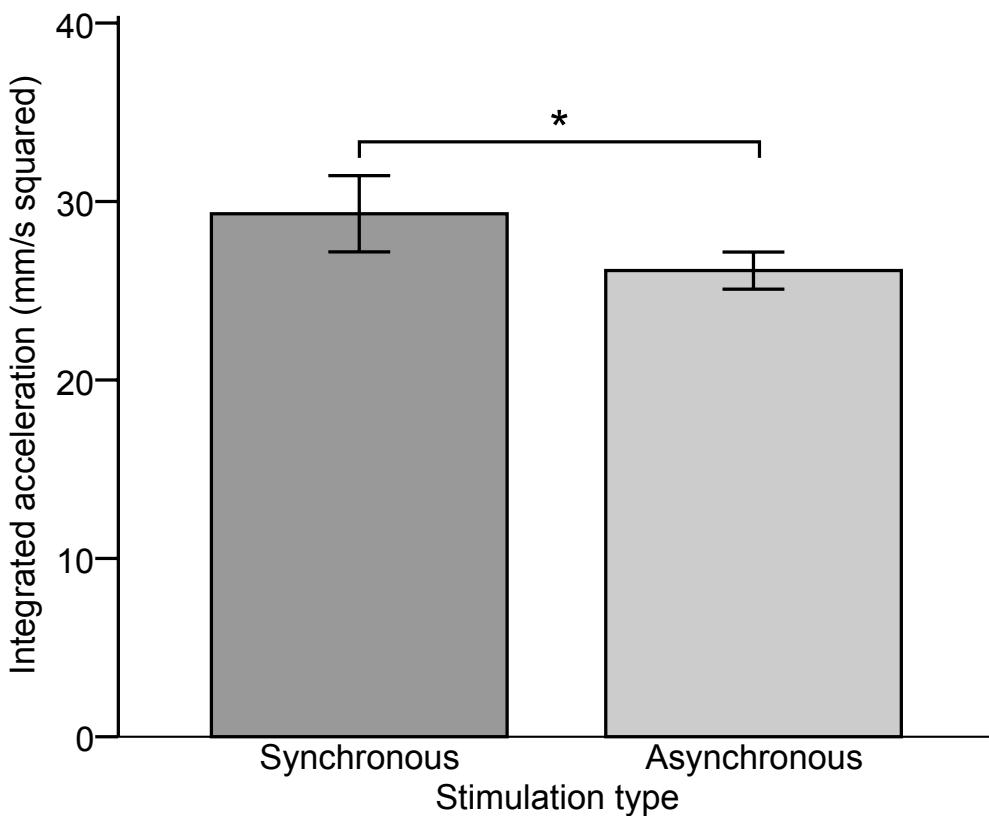
A 2x2x3 mixed ANOVA was performed for integrated acceleration with Group (Low AQ vs. High AQ) as a between-subjects factor and Stimulation Type (Synchronous vs. Asynchronous) and Rubber Arm Position (Position 1 vs. Position 2 vs. Position 3) as within-subjects factors.

A significant main effect of Group was found for integrated acceleration,  $F(1, 21) = 8.19$ ,  $p < .01$ , Cohen's  $d = 1.19$ . The low AQ group displayed greater integrated acceleration across conditions ( $M = 31.36$ ,  $SD = 7.66$ ) than the high AQ group ( $M = 23.84$ ,  $SD = 4.63$ ; *Figure 4*). Consistent with our previous examination of the acceleration of reaching movements in the rubber-hand illusion (Paton et al.,

2012), a significant main effect of Stimulation Type was found for integrated acceleration,  $F(1, 21) = 6.19, p < .05$ , Cohen's  $d = 0.47$ . However, contrary to the direction of our previous finding for nonclinical participants, synchronous stimulation ( $M = 29.70, SD = 10.50$ ) was associated with greater integrated acceleration in subsequently performed reach-to-grasp movements than asynchronous stimulation ( $M = 25.82, SD = 4.99$ ; *Figure 5*). This inconsistency between studies may reflect the difference in the acceleration variables used previously (recorded via an accelerometer) and in the current study (recorded via a 6-dimensional tracker, and derived specifically from the linear axes, thus controlling for rotational changes that occur throughout the movement). This previous study also contained a number of independent variables not included in the present study, potentially contributing to a difference between studies in the stimulation-type comparison. No other main or interaction effects were found for this variable ( $p > .05$ ).



**Figure 4.** Integrated acceleration of reach-to-grasp movements between participant groups separated by AQ scores. Error bars indicate  $\pm 1$  standard error. (\*\* $p < .01$ ).

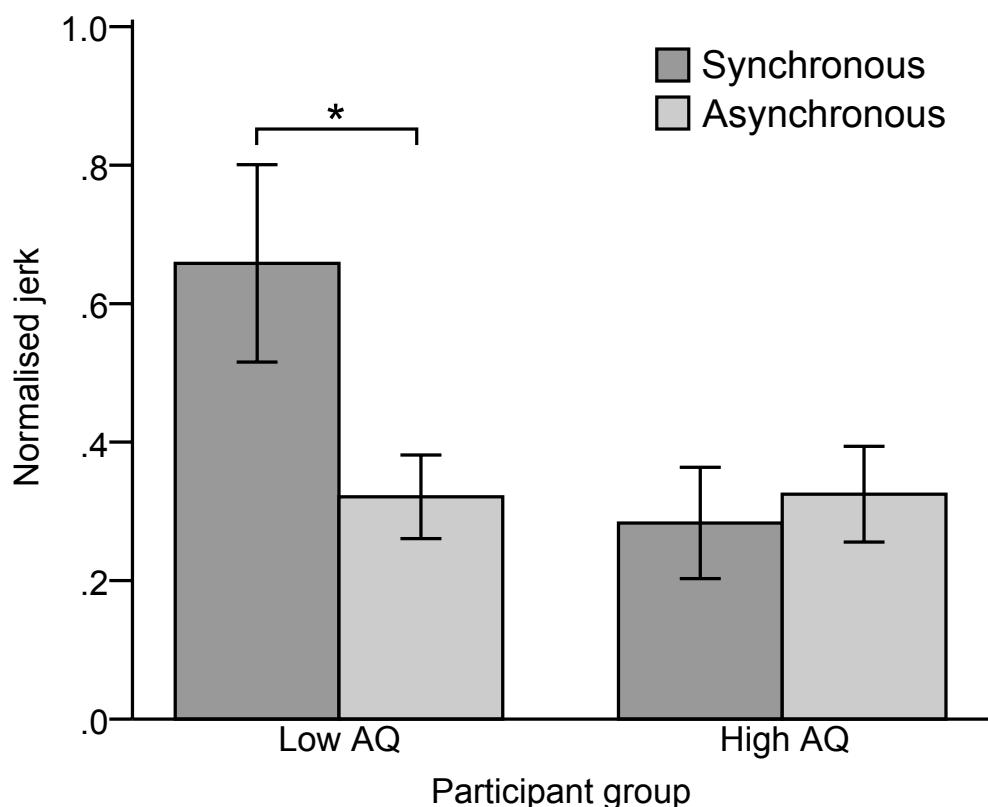


**Figure 5.** Integrated acceleration of reach-to-grasp movements between stimulation types. Error bars indicate  $\pm 1$  standard error. (\* $p < .05$ ).

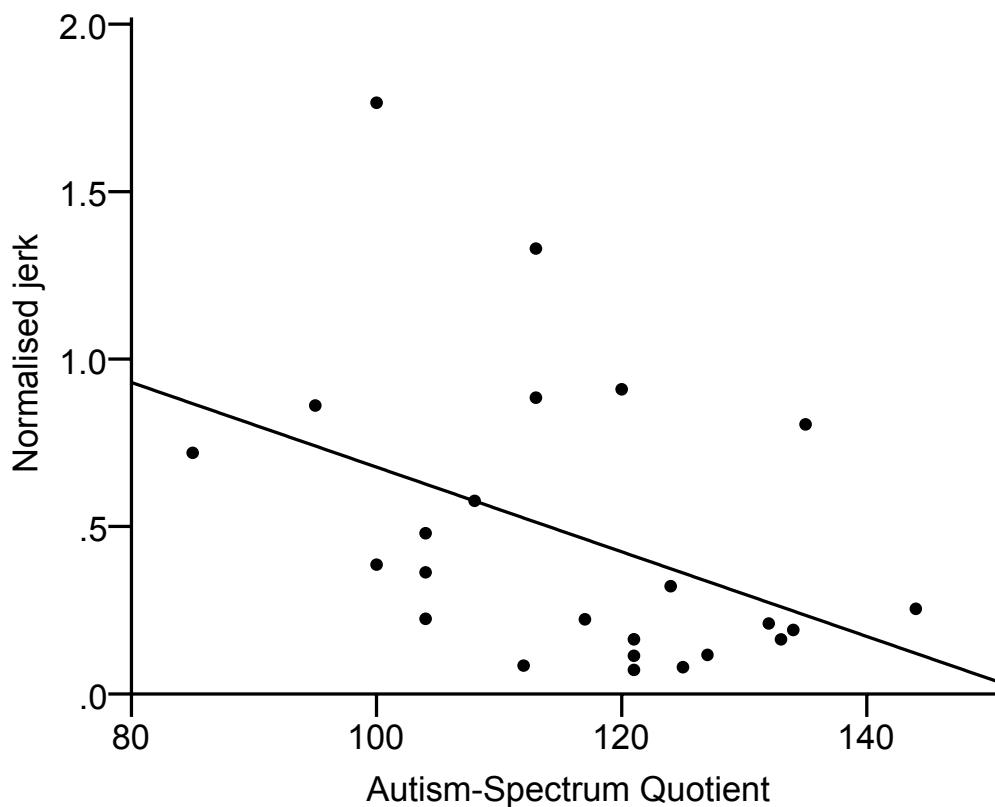
### 3.4 Normalised integrated jerk

Differences in integrated jerk between the two stimulation conditions were examined separately for each AQ group using the nonparametric Wilcoxon signed-rank test. For the low AQ group, the integrated jerk of reaching movements was significantly greater following synchronous stimulation ( $M = .66$ ,  $SD = .49$ ) than when following asynchronous stimulation ( $M = .32$ ,  $SD = .21$ ),  $z = -2.28$ ,  $p < .05$ , Cohen's  $d = 0.89$ . In contrast, integrated jerk did not differ significantly between stimulation conditions for the high AQ group,  $z = -0.24$ ,  $p = .81$ , Cohen's  $d = -0.16$  (*Figure 6*). Similarly, post-hoc correlational analyses conducted to further elucidate this effect indicated that AQ scores shared a significant negative correlation with integrated jerk following synchronous stimulation ( $r_s = -.47$ ,  $p < .05$ , two-tailed),

but not asynchronous stimulation ( $r_s = -.11$ ,  $p = .63$ , two-tailed; *Figure 7*). These findings therefore indicate that the integrated jerk of reach-to-grasp movements was increased by the presence of the illusion for participants who scored lower on the AQ, but was not significantly different across stimulation types for participants who scored higher on the AQ.



**Figure 6.** Normalised integrated jerk of reach-to-grasp movements across AQ groups and stimulation types. Error bars indicate  $\pm 1$  standard error. (\* $p < .05$ ).



**Figure 7.** A significant correlation was observed between AQ scores and the normalised integrated jerk of reach-to-grasp movements following synchronous stimulation. ( $r_s = -.47, p < .05$ , two-tailed; linear least squares regression line of best fit:  $y = -0.013x + 1.941, R^2 = .18, t = -2.17, p < .05$ ).

### 3.5 Movement duration

For movement duration, a  $2 \times 2 \times 3$  mixed ANOVA was performed with Group (Low AQ vs. High AQ) as a between-subjects factor and Stimulation Type (Synchronous vs. Asynchronous) and Rubber Arm Position (Position 1 vs. Position 2 vs. Position 3) as within-subjects factors. A significant two-way interaction between Stimulation Type and Group was found,  $F(1, 22) = 4.48, p < .05$ . No other main or interaction effects were found for this variable ( $p > .05$ ).

Post hoc tests comparing movement duration between stimulation conditions did not reach significance, however, for either the low AQ group (synchronous

stimulation:  $M = 1.71$ ,  $SD = 0.37$ ; asynchronous stimulation:  $M = 1.53$ ,  $SD = .40$ ;  $t(11) = 2.14$ ,  $p = .06$ , Cohen's  $d = 0.46$ ) or the high AQ group (synchronous stimulation:  $M = 1.55$ ,  $SD = 0.50$ ; asynchronous stimulation:  $M = 1.62$ ,  $SD = .41$ ;  $t(11) = -0.94$ ,  $p = .37$ , Cohen's  $d = -0.13$ ).

### 3.6 Further reach measures

For each of the remaining reach measures, a 2x2x3 mixed ANOVA was performed with Group (Low AQ vs. High AQ) as a between-subjects factor and Stimulation Type (Synchronous vs. Asynchronous) and Rubber Arm Position (Position 1 vs. Position 2 vs. Position 3) as within-subjects factors. No significant differences were found across conditions or groups for mean velocity, peak velocity, or maximum horizontal displacement ( $p > .05$ ). Significant differences observed for relative time to peak velocity and angle of initial movement are reported in Supplementary Material. Means and standard deviations for each reach measure are shown in *Table S1* in Supplementary Material.

## 4. Discussion

The present study supports the hypothesis that proprioceptive and sensorimotor characteristics of ASD, as reflected in the multimodal effects of the rubber-hand illusion, vary together with autistic traits in the general population. Nonclinical adults scoring higher on autistic traits showed reduced effects of the illusion on perceived arm position compared to those scoring lower on autistic traits (as indicated by a lesser influence of the position of the rubber arm on estimated arm position during synchronous stimulation). Individuals with higher autistic traits also demonstrated reduced sensitivity to the presence of the illusion in their reaching movements. These effects occurred despite both groups reporting the typical subjective effects of the illusion, concerning referral of touch and a heightened sense of ownership for the rubber hand. This pattern of intact subjective effects but diminished proprioceptive and sensorimotor effects resembles that found previously for the rubber-hand illusion in a sample of adults

diagnosed with ASD (Paton et al., 2012). The present findings are also consistent with a study of the rubber-hand illusion in children diagnosed with ASD, which reports intact subjective effects of the illusion but delayed proprioceptive effects (Cascio et al., 2012). While autistic traits are most commonly defined in terms of social difficulties and unusual repetitive behaviours and interests, the present study adds to recent research that has found sensory differences associated with ASD to vary together with other aspects of this condition in the general population (e.g., Donohue et al., 2012; Grinter, Maybery, et al., 2009; Grinter, Van Beek, et al., 2009; Walter et al., 2009). This research is therefore consistent with the continuum hypothesis of ASD (Happé et al., 2006; Mandy & Skuse, 2008), and highlights the relevance of sensory characteristics in defining a broader autistic phenotype.

To characterise the proprioceptive and sensorimotor differences associated with autistic traits in the present study, we need to emphasise a distinction between different levels of sensory integration in the rubber-hand illusion. Visuotactile integration in the illusion is dependent upon the close temporal synchrony of repetitive tactile and visual inputs (Botvinick & Cohen, 1998; Shimada et al., 2009), a signal that is conveyed precisely during illusion-induction by continuous tactile and visual stimulation. In contrast, we can hypothesise that changes in perceived arm position induced by the illusion reflect integration between sensory (proprioceptive) estimates of arm position and predictions for arm position derived from the illusory context. This context would, for example, include the visual presence of the rubber arm and the (illusory) location of felt touch. The evidence that proprioceptive differences occurred despite a typical subjective experience of the illusion in each group suggests that enhanced proprioceptive performance in the high AQ group reflects increased reliance on sensory (proprioceptive) input at the expense of the more global context. This distinction between visuotactile and proprioceptive mechanisms is consistent with a model of the rubber-hand illusion proposed by Makin et al. (2008), in which changes in perceived arm position occur subsequent to visual capture of the tactile input, based on evidence disassociating the co-dependence, time course, and spatial extent of these effects.

The enhanced proprioceptive performance of individuals with autistic traits in the present study (and in the ASD group in our previous rubber-hand illusion study; Paton et al., 2012) conflicts somewhat with a recent study examining limb proprioception in ASD outside of the context of sensory illusions (Fuentes, Mostofsky, & Bastian, 2011). In particular, no differences were found in this latter study in the accuracy and precision of proprioceptive estimates regarding arm position between individuals with ASD and healthy controls. As described in the preceding paragraph, enhanced performance in the present study can be explained in terms of a reduced tendency for taking the wider context, here provided by the rubber-hand illusion, into account. The more accurate proprioceptive performance in the high AQ group may then be serendipitous given the specific (illusion-based) task context used, rather than reflecting a superior capacity for accurate proprioception in ASD. The apparent conflict between the present findings and those of Fuentes et al. (2011) may therefore reflect a lack of a modulating context in the experimental set-ups used in the latter study, such that the hypothesised disregard for contextually-based models of sensory input in autistic perception did not cause a deviation in the performance of individuals with ASD from controls. The implication here is that group differences in performance may vary across tasks depending on whether there is a more global model of sensory input suggested by the specific task context, and whether this task-specific context aids or misleads accurate performance. Integrating proprioceptive sensory information into a more global model (as we suggest occurs for the low AQ group in the rubber-hand illusion) is likely to be beneficial to accuracy in some contexts but not others.

Within a predictive processing framework of perception, the degree of precision that is expected from the sensory input in a given context determines the relative contribution of (top-down) relatively global hypotheses regarding the state of the world and (bottom-up) sensory input (Feldman & Friston, 2010; Friston & Stephan, 2007). The former are more likely to mediate an influence of contextual information on perception. We can therefore speculate that the group differences observed in the present study reflect individual variation in the expected levels of sensory precision (Friston et al., 2013; Hohwy, 2013; Paton et al., 2012). In particular, the greater tendency of the low AQ group to draw on the illusory

context to estimate arm position, as suggested by greater proprioceptive drift when the distance between the real and rubber arms was increased in the synchronous stimulation condition, may reflect an expectation for low precision of bottom-up sensory (proprioceptive) estimates within the unusual context of the illusion. Similarly, the enhanced proprioceptive performance of the high AQ group may reflect an expectation for high precision in sensory input, leading to a lessened influence of global models that take into account the illusory context when estimating arm position. This would explain why these participants specifically show diminished proprioceptive effects of the illusion rather than reduced effects of the illusion in general. The latter could otherwise be explained by a general bias concerning global integration.

A group difference was also found in the effect of the rubber-hand illusion on the smoothness of reach-to-grasp movements performed with the stimulated hand following illusion-induction. In particular, we found that individuals lower in autistic traits executed movements less smoothly following synchronous stimulation (as indicated by increased normalised integrated jerk) compared to the asynchronous stimulation control condition. The high AQ group, in contrast, showed uniform movement smoothness across conditions, at a level similar to that of the asynchronous condition for the low AQ group. This is partially consistent with our previous study (Paton et al., 2012), which observed differences in the integrated acceleration of reach-to-grasp movements following the rubber-hand illusion between a clinical ASD group and a nonclinical control group. This previous clinical study did not assess movement smoothness, however, and there was no group difference in integrated acceleration in the present study involving nonclinical participants – so a direct comparison between reaching movements found for clinical ASD and nonclinical autistic traits is a task for further studies.

The reach effects in the present study can also be interpreted in terms of group differences in expectations for sensory precision, which we brought to bear on the proprioceptive drift findings above. Again, given that individuals with higher autistic traits report experiencing the typical subjective effects of the illusion, the lack of difference in reaching movements across stimulation conditions seems best

explained by insensitivity to the context of the illusion when executing movement rather than a general resistance to the illusion itself. Explicating this notion of context-insensitivity within the framework of predictive processing, we can hypothesise that a less smooth movement would be performed when the individual expects imprecision in their proprioceptive and kinaesthetic feedback for the planned movement. This could occur due to difficulty coordinating movement when the trajectory required to reach the target is uncertain, or, similarly, could reflect the introduction of exploratory movements to elicit proprioceptive and kinaesthetic feedback. In contrast, a more confident, or smoothly executed, movement may be likely to occur when the individual assumes high precision in their estimate of initial arm position and predicts high precision in their proprioceptive and kinaesthetic feedback once the movement is underway. Quantifying this in terms of differences of higher order temporal derivatives (e.g., jerk) is useful because one may assume that since such derivatives encompass relatively long time-scales they are encoded at higher cortical levels, consistent with the idea of more high-level, relatively global context modulation (in essence, trying to anticipate the overall smoothness of the movement given levels of expected uncertainty; cf. Friston et al., 2010).

An alternative explanation is that the observed differences in reaching movements for the low AQ group are directly related to group differences in the magnitude of drift in arm location induced by the illusion. Specifically, the reduced smoothness of movement following the synchronous stimulation condition for the low AQ group could reflect the increased tendency for proprioceptive drift that this group demonstrates. Counter to this interpretation, however, is the lack of difference found in the displacement measures of the reach-to-grasp movements. If the shift in perceived arm location towards the rubber arm contributed significantly to the subsequent reaching movements, we would expect this to manifest as a difference in the angle of initial movement or peak horizontal displacement of the reach trajectories. For example, a shift in perceived arm location to the left would mean that the subsequent arm trajectory would have a sharper angle of movement to the right and greater deviation to the right than that really required to reach the target. This logic is adopted in Newport et al. (2010), Zopf et al. (2011), and Heed

et al. (2011), who each report differences in reach displacement induced by the rubber-hand illusion (see also Kammers, de Vignemont, et al., 2009, who report no differences in the displacement of reaching movements following the illusion; Kammers, Longo, Tsakiris, Dijkerman, & Haggard, 2009; Kammers, Verhagen, et al., 2009). Given that we did not see differences in displacement parameters in the present study, the observed effect of the illusion on movement smoothness may not merely reflect the increased proprioceptive drift experienced by the low AQ group.

The interpretation of the reach data that we favour leads to an interesting implication regarding movement impairments that commonly occur in ASD (e.g., Mari et al., 2003; Nazarali et al., 2009; Rinehart et al., 2001; Rinehart, Tonge, et al., 2006) (see Fournier, Hass, Naik, Lodha, & Cauraugh, 2010, for meta-analysis). If, as we suggest in the preceding paragraphs, the lack of differences in reaching movements for the high AQ group across synchronous and asynchronous stimulation conditions reflects insensitivity to the context-specific inducements of expectations for imprecision in the sensory estimates used to guide reaching movements, then a tendency to disregard context in this manner in ASD may lead to overconfident movement in contexts that would usually suggest imprecision. Difficulties in movement coordination might then be partly explained in terms of movement errors caused by overconfident movement execution, and would be specifically expected to occur in contexts that advise for tentative movement execution. In other words, we suggest that a lesser sensitivity to, and urge to resolve, ambiguity in body position may contribute to uncoordinated movement in ASD. It might therefore be useful to further examine the effects of the rubber-hand illusion on movement execution with respect to clinical measures of motor coordination in ASD.

In summary, the present study examined individual differences in the relative contribution of sensory input and contextual factors to perception. Working within a Bayesian (prediction error minimisation) framework, we reasoned that autistic sensory integration involves a tendency to ignore contextual information that suggest imprecision, and predicted that this entails high estimations of sensory

precision across contexts, leading to an increased reliance on lower-level sensory estimates and a decreased tendency to subsume input under higher-level expectations. The finding that, following synchronous stimulation, nonclinical individuals high in autistic traits show reduced sensitivity to the position of the rubber hand in their proprioceptive estimates and show less sensitivity to uncertainty while executing reaching movements, despite reporting the subjective experience of the illusion, is consistent with this hypothesis. The ability to modulate expected levels of sensory precision in response to contextual information suggesting varying uncertainty may lead to a better understanding of the complex constellation of compromised and enhanced perceptual performance in ASD, as well as of individual differences in perception in the general population.

## Acknowledgements

The authors wish to thank Uta Frith for very helpful discussions regarding study design and theoretical issues. This work was supported by an Australian Research Council Discovery grant (DP1311336). JH is supported by an Australian Research Council Future Fellowship (FT100100322). PE is supported by a NHMRC Clinical Research Fellowship (546244). The authors declare no conflicts of interest.

## S1. Supplementary Results

### S1.1 Further post-hoc analyses for illusion ratings

As described in Section 3.1. of the main text, an interaction effect was found between Stimulation Type and Group,  $F(1, 22) = 4.53, p < .05$ . Post-hoc tests indicated that the low AQ group rated questionnaire items higher following synchronous stimulation ( $M = 10.59, SD = 2.59$ ) than following asynchronous stimulation ( $M = 5.14, SD = 3.19$ ),  $t(11) = 6.61, p < .001$ , Cohen's  $d = 1.88$ . The high AQ group also rated items higher following synchronous stimulation ( $M = 11.01, SD = 2.46$ ) compared to asynchronous stimulation ( $M = 7.66, SD = 2.96$ ),  $t(11) = 6.21, p < .001$ , Cohen's  $d = 1.23$ . Further post-hoc tests indicated that groups didn't differ in their item ratings across the questionnaire following synchronous stimulation,  $t(22) = -0.41, p = .69$ , Cohen's  $d = -0.17$ , or following asynchronous stimulation,  $t(22) = -2.01, p = .06$ , Cohen's  $d = -0.82$ .

There was also a three-way interaction effect between Item Type, Stimulation Type, and Group,  $F(1, 22) = 6.64, p < .05$ . To clarify this interaction effect, two-way repeated measures ANOVAs were conducted separately for each AQ group, with Item Type and Stimulation Type as factors. There was a significant main effect of Item Type for both the low AQ group,  $F(1, 11) = 21.08, p < .01$ , Cohen's  $d = 1.29$ , and the high AQ group,  $F(1, 11) = 14.35, p < .01$ , Cohen's  $d = 0.69$ . This indicates that both groups rated illusion items (low AQ group:  $M = 9.73, SD = 3.12$ ; high AQ group:  $M = 10.27, SD = 2.58$ ) higher than control items (low AQ group:  $M = 6.00, SD = 2.65$ ; high AQ group:  $M = 8.40, SD = 2.80$ ). There was a significant main effect of Stimulation Type for both the low AQ group,  $F(1, 11) = 43.74, p < .001$ , Cohen's  $d = 1.87$ , and the high AQ group,  $F(1, 11) = 38.51, p < .001$ , Cohen's  $d = 1.23$ . This indicates that both groups made higher item ratings following synchronous stimulation (low AQ group:  $M = 10.58, SD = 2.59$ ; high AQ group:  $M = 11.01, SD = 2.46$ ) compared to asynchronous stimulation (low AQ group:  $M = 5.14, SD = 3.19$ ; high AQ group:  $M = 7.66, SD = 2.96$ ). There was a significant Item Type by Stimulation Type interaction effect for both the low AQ group,  $F(1, 11) = 49.42, p < .001$ , and the high AQ group,  $F(1, 11) = 27.51, p < .001$ . Post-hoc tests for the low

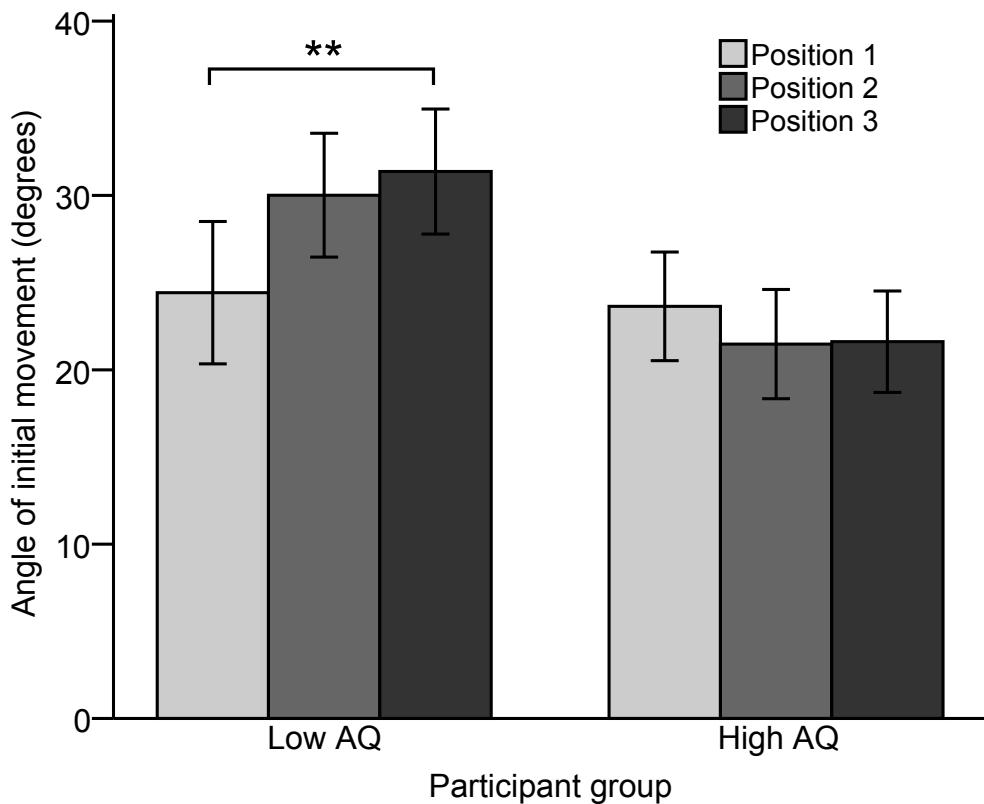
AQ group indicated that synchronous stimulation was associated with significantly higher ratings on illusion items ( $M = 14.52, SD = 3.35$ ) than control items ( $M = 6.65, SD = 3.10$ ),  $t(11) = 7.05, p < .001$ , Cohen's  $d = 2.44$ , while there was no significant difference between illusion and control items for asynchronous stimulation,  $t(11) = -0.45, p = .66$ , Cohen's  $d = -0.11$ . Similarly, post-hoc tests for the high AQ group indicated that synchronous stimulation was associated with significantly higher ratings on illusion items ( $M = 13.07, SD = 2.45$ ) than control items ( $M = 8.94, SD = 2.93$ ),  $t(11) = 6.39, p < .001$ , Cohen's  $d = 1.53$ , while there was no significant difference between illusion and control items for asynchronous stimulation,  $t(11) = -0.57, p = .58$ , Cohen's  $d = -0.12$ . Further post-hoc tests indicated that illusion items were rated higher during synchronous stimulation (low AQ group:  $M = 14.52, SD = 3.35$ ; high AQ group:  $M = 13.07, SD = 2.45$ ) than asynchronous stimulation (low AQ group:  $M = 4.95, SD = 4.19$ ; high AQ group:  $M = 7.47, SD = 3.52$ ) for both the low AQ group,  $t(11) = 7.66, p < .001$ , Cohen's  $d = 2.52$ , and the high AQ group,  $t(11) = 6.11, p < .001$ , Cohen's  $d = 1.85$ . Differences in control item ratings across synchronous (low AQ group:  $M = 6.65, SD = 3.10$ ; high AQ group:  $M = 8.94, SD = 2.93$ ) and asynchronous (low AQ group:  $M = 5.34, SD = 2.71$ ; high AQ group:  $M = 7.85, SD = 2.79$ ) stimulation approached significance for the high AQ group but did not survive Bonferroni correction of significance levels,  $t(11) = 3.28, p = .007$  (corrected  $\alpha = .00625$ ), Cohen's  $d = 0.38$ , nor was this comparison significant for the low AQ group,  $t(11) = 1.89, p = .09$ , Cohen's  $d = 0.45$ .

## S1.2 Relative time to peak velocity

There was a significant main effect of Rubber Arm Position for the relative time to peak velocity measure,  $F(2, 44) = 5.60, p < .01$ . No other main or interaction effects were found for this variable ( $p > .05$ ). Post hoc tests indicated that it took significantly longer to reach peak velocity for position 2 (25 cm separation between the participant's arm and the rubber arm;  $M = 33.51, SD = 8.12$ ) compared to position 1 (20 cm separation;  $M = 28.26, SD = 8.46$ ),  $p < .01$ , Cohen's  $d = -0.63$ .

### S1.3 Angle of initial movement

For the angle of initial movement, Mauchly's Test of Sphericity indicated that the assumption of sphericity was violated ( $p < .05$ ), so Greenhouse-Geisser correction was used for this variable. A significant interaction existed between Rubber Arm Position and Group,  $F(2, 35) = 4.61, p < .05$ . To clarify this effect, one-way ANOVAs were performed separately for each AQ group with Rubber Arm Position as the factor. These tests indicated that the low AQ group showed a significant difference between rubber arm positions,  $F(1, 14) = 5.22, p < .05$ , while the high AQ group did not,  $F(1, 14) = 0.58, p = .50$ . Follow up post hoc tests revealed that the low AQ group had a significantly greater angle of initial movement for position 3 (30 cm separation between the participant's arm and the rubber arm;  $M = 31.37, SD = 12.42$ ) than for position 1 (20 cm separation;  $M = 24.42, SD = 14.14$ ),  $p < .01$ , Cohen's  $d = 0.52$  (*Figure S1*). These values indicate that for the low AQ group the initial hand movement during the reaching phase was on average directed further laterally when the rubber arm was placed in position 3, and more directly forward for position 1. This effect was not specific to synchronous stimulation, however. No other main or interaction effects were found for this variable ( $p > .05$ ).



**Figure S1.** Angle of initial movement across rubber arm position conditions between participant groups separated by AQ scores. Error bars indicate  $\pm 1$  standard error.

(\*\* $p < .01$ ).

**Table S1.**

Means and standard deviations of reach measures across conditions separated by Stimulation Type and Rubber Arm Position.

Measure	Unit	Synchronous stimulation			Asynchronous stimulation		
		Position 1	Position 2	Position 3	Position 1	Position 2	Position 3
Movement duration	s	1.71 (0.54)	1.55 (0.47)	1.62 (0.47)	1.66 (0.56)	1.57 (0.41)	1.48 (0.45)
Mean velocity	mm/s	173.06 (49.87)	175.48 (49.46)	177.78 (43.64)	176.61 (52.98)	178.03 (39.82)	185.70 (47.28)
Peak velocity	mm/s	417.87 (135.70)	390.97 (131.49)	427.00 (129.14)	403.37 (116.61)	407.11 (125.84)	419.38 (126.26)
Relative time to peak velocity	%	27.97 (8.61)	34.65 (10.46)	30.21 (11.24)	28.56 (12.18)	32.38 (9.23)	31.04 (8.78)
Peak horizontal displacement	mm	56.94 (23.24)	54.53 (21.23)	57.90 (20.71)	53.03 (23.27)	53.92 (21.16)	53.36 (20.77)
Angle of initial movement	°	24.72 (12.45)	24.77 (14.23)	28.26 (12.96)	23.34 (13.52)	26.72 (13.26)	24.72 (13.67)
Integrated acceleration	mm/s <sup>2</sup>	28.90 (9.44)	27.81 (9.51)	32.26 (17.55)	25.67 (5.01)	26.77 (7.19)	25.09 (5.08)
Normalised jerk	-	0.57 (0.82)	0.37 (0.43)	0.47 (0.55)	0.39 (0.40)	0.32 (0.25)	0.26 (0.23)

## Linking text between chapters 2 and 3

In *Chapter 2*, susceptibility to the rubber-hand illusion was examined in a sample of nonclinical adults. Proprioceptive and motoric effects of the illusion were found to differ in this sample with respect to variation in autism-like traits. These findings are partially replicated and extended in *Chapter 3*, which reports data from adults with formal ASD diagnoses in addition to nonclinical individuals that vary in their autism-like traits. *Chapter 3* also characterises the motoric effects of the illusion in greater detail; we find that it is not the *initiation* of movement that is influenced by the illusion, but rather *later* stages of movement, consistent with a difference in ASD in the integration between sensory signals received during movement and expectations for arm position influenced by the illusion. These findings are discussed in the context of predictive processing theories of neurocognitive function. (For further background to this chapter, please refer back to the linking text between *Chapter 1* and *Chapter 2*.)

## Declaration for Thesis Chapter 3

### Declaration by candidate

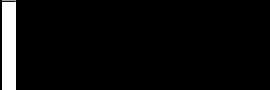
In the case of *Chapter 3*, the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
<ul style="list-style-type: none"> <li>Contributed to the experiment design.</li> <li>Performed data analysis.</li> <li>Contributed to the interpretation of results.</li> <li>Wrote the paper.</li> </ul>	50%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Bryan Paton	Technical input to data collection and analysis. Contributed to the experiment design, interpretation of results, and writing.	10%
Melissa Kirkovski	Data collection.	15%
Peter Enticott	Contributed to the experiment design, interpretation of results, and writing.	N/A
Jakob Hohwy	Contributed to the experiment design, interpretation of results, and writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date 16/02/2016
--------------------------	---	--------------------

Main Supervisor's Signature		Date 16/02/2016
-----------------------------------	---	--------------------

# Chapter 3

## Context sensitivity in action decreases along the autism spectrum: A predictive processing perspective

Colin J. Palmer<sup>1</sup>, Bryan Paton<sup>1, 2, 3</sup>, Melissa Kirkovski<sup>4, 5</sup>, Peter G. Enticott<sup>4, 5</sup> and Jakob Hohwy<sup>1</sup>

<sup>1</sup>*Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup>*School of Psychological Sciences, Monash University, Clayton, VIC 3800, Australia*

<sup>3</sup>*Monash Biomedical Imaging, Monash University, Clayton, VIC 3168, Australia*

<sup>4</sup>*Monash Alfred Psychiatry Research Centre, The Alfred and Central Clinical School, Monash University, Melbourne, VIC 3004, Australia*

<sup>5</sup>*Cognitive Neuroscience Unit, School of Psychology, Deakin University, Burwood, VIC 3125, Australia*

Published in *Proceedings of the Royal Society B: Biological Sciences*, 2015

## Abstract

Recent predictive processing accounts of perception and action point towards a key challenge for the nervous system in dynamically optimising the balance between incoming sensory information and existing expectations regarding the state of the environment. Here we report differences in the influence of the preceding sensory context on motor function, varying with respect to both clinical and subclinical features of autism spectrum disorder (ASD). Reach-to-grasp movements were recorded subsequent to an inactive period in which illusory ownership of a prosthetic limb was induced. We analysed the sub-components of reach trajectories derived using a minimum-jerk fitting procedure. Nonclinical adults low in autistic features showed disrupted movement execution following the illusion compared to a control condition. In contrast, individuals higher in autistic features (both those with ASD and nonclinical individuals high in autistic traits) showed reduced sensitivity to the presence of the illusion in their reaching movements while still exhibiting the typical perceptual effects of the illusion. Clinical individuals were distinct from nonclinical individuals scoring high in autistic features, however, in the early stages of movement. These results suggest that the influence of high-level representations of the environment differs between individuals, contributing to clinical and subclinical differences in motor performance that manifest in a contextual manner. Since high-level representations of context help to explain fluctuations in sensory input over relatively longer time scales, more circumscribed sensitivity to prior or contextual information in autistic sensory processing could contribute more generally to reduced social comprehension, sensory impairments and a stronger desire for predictability and routine.

## 1. Introduction

An influential idea in cognitive science is that for the brain to successfully represent and interact with its environment it engages in an unconscious process of inference about the external causes of sensory stimulation. This idea arises in response to the ambiguous relationship between sensory inputs and worldly states, which seems to necessitate that sensory information is integrated with prior and contextual information regarding the likely causes of input. Contemporary *predictive processing* accounts of cortical function provide a computationally and biologically plausible mechanism through which this process might occur via the implementation of probabilistic generative models (Friston, 2005, 2009). In a recurrent hierarchical arrangement, hypotheses regarding the present causes of input are used to generate predictions of sensory activity at subordinate levels. Hypotheses at each level are then updated iteratively to more closely match predictions to incoming data. In this manner, a dynamic representation of the causal structure of the world comes to be encoded across the neocortex, graded from lower to higher levels of spatial and temporal abstraction. Action is situated within this framework as a process of manipulating the sensory input to match predictions (e.g., predictions regarding the parameters of unfolding proprioceptive feedback) (Adams, Shipp, et al., 2013; Shipp et al., 2013).

When the brain is cast in this light, we gain a nuanced perspective on how systematic differences in perception and action may emerge between individuals. To improve predictions over time, the influence of sensory input on cortical representations must be weighted by how informative the input is expected to be concerning regularities in the world (i.e., weighted in proportion to the expected precision of the input relative to the precision of existing expectations) (Feldman & Friston, 2010). This captures the intuitive principle that sensory information should be drawn upon to a greater degree (at the expense of prior or contextual information) in contexts when the present input is expected to be more highly reliable in determining the state of the external world. Thus, a key task for the nervous system is in optimising the relative influence of (top-down) prior or contextual information on low level, local processing. This task can be challenging

because different contexts require recruitment of different levels of the cortical hierarchy to accurately represent the causal structure of the world. Changes in sensory input could be best accounted for by inferring the presence of either shorter or longer term regularities, for example, and could reflect changes in first- or second-order statistics (for discussion, see Hohwy, 2013).

These concepts have been drawn upon very recently to understand ASD and nonclinical variance in autistic features (Brock, 2012; Hohwy, 2013; Lawson et al., 2014; Palmer et al., 2013; Pellicano & Burr, 2012; Skewes et al., 2014; van Boxtel & Lu, 2013b; Van de Cruys et al., 2014). ASD is a highly prevalent developmental condition (~1%) characterised in significant part by social, communicative and behavioural atypicalities (American Psychiatric Association, 2013; Lai et al., 2014). Other well-established features include sensory hyper- and hypo-sensitivities, a detail-oriented processing style, a strong preference for predictability and routine, cognitive inflexibility and poor motor coordination (Fournier et al., 2010; Gowen & Hamilton, 2013; Happé & Frith, 2006; Simmons et al., 2009). Social and non-social autistic characteristics vary to a significant degree across the general population, in both children and adults (e.g., Baron-Cohen et al., 2001; Constantino & Todd, 2003; Hurst et al., 2007; Posserud et al., 2006). Pellicano and Burr (2012) argue that nonsocial features of ASD can be understood as a reduced influence of prior experience on sensory processing (see also Mitchell & Ropar, 2004). Within the predictive processing framework, this idea has been developed in terms of an increased effect of sensory stimulation on cortical representations of the world such that perception is bound more closely to lower levels of representation – where the weighting of sensory input is tied to estimations of state-dependent uncertainty (Hohwy, 2013; Lawson et al., 2014; Van de Cruys et al., 2014). As mentioned, differences between individuals in the modulation of lower-level processing by higher-level representations can arise because there is not an unequivocal answer regarding the appropriate levels of the hierarchy to recruit in a given situation. Work has just begun in unpacking the implications of this type of account for our understanding of sensory, motor and social symptoms in ASD, as well as for individual differences more broadly (e.g., Gómez et al., 2014; Hohwy &

Palmer, 2014; Lawson et al., 2014; Palmer et al., 2013; Skewes et al., 2014; van Boxtel & Lu, 2013a; Van de Cruys et al., 2014).

Here we examine the influence of the preceding sensory context on sensorimotor function with respect to both clinical and subclinical features of ASD. Specifically, we investigate how reach-to-grasp movements unfold following exposure to the rubber-hand illusion (RHI) – a multisensory illusion of ownership for a prosthetic limb (Botvinick & Cohen, 1998; Ehrsson, 2012). Induction of this illusion (via synchronous tactile stimulation of a visible prosthetic limb and the occluded real limb) influences bodily representations for perception and action, reflected, for example, in drift in perceived arm position towards the prosthetic limb (Tsakiris & Haggard, 2005; Zopf et al., 2011). Moving the arm subsequent to an inactive period of illusion induction is therefore likely to require integration between prior, context-sensitive expectations regarding limb position and sensory (proprioceptive) feedback received once movement is underway. This provides a novel setting for examining consequences of the dynamic interaction between sensory evidence and higher-order expectations specified by predictive processing. In this type of paradigm, sensitivity to the context of the illusion can be understood in terms of the relative influence of higher-level representations filtering down the cortical hierarchy to modulate predictions at lower levels. We therefore expected that divergence between individuals across illusory and non-illusory conditions would be revealing in terms of the processing imbalances hypothesised by inferential accounts of ASD. In particular, we expect individuals higher in autistic features to be increasingly disinclined to let higher-level representations be informative about low-level sensory input; conversely, they should be more inclined to consider their sensory input informative.

Nonclinical individuals grouped by their level of autistic traits show differences in sensitivity to the presence of the RHI in reaching movements (Palmer et al., 2013) (see also Paton et al., 2012). Specifically, individuals low in autistic features exhibit reduced smoothness of movement following the illusion compared to a control condition, while individuals high in autistic features show uniformly smooth movements across conditions. In the present study, we examined whether adults

with ASD demonstrate similarly reduced sensitivity in movement to the preceding context of the RHI. Crucially, we compare those with a diagnosis of ASD with both nonclinical individuals high in autistic features and nonclinical individuals low in autistic features. This allows us to assess how the predictive processing account of individual differences we described and developed above coheres with clinical and nonclinical features of movement. Additionally, we decompose reaching trajectories into sub-components (described in Methods) to more closely examine whether differences in action following the illusion are consistent with differences in the context-sensitive integration of sensory feedback with prior expectations.

## 2. Methods

### 2.1 Participants

Three participant groups were involved in this experiment. Thirty nonclinical adults were recruited via university and hospital advertisements and separated into two groups based on a median-split of Autism-Spectrum Quotient scores (AQ, an adult inventory measure of social and non-social autistic traits; Baron-Cohen et al., 2001). Thus, we examined a *Low AQ* group of 15 nonclinical individuals (8 female; age:  $M = 30.20$ ,  $SD = 7.31$  years; AQ:  $M = 8.07$ ,  $SD = 3.96$ ) and a *High AQ* group of 15 nonclinical individuals (5 female; age:  $M = 29.87$ ,  $SD = 8.61$  years; AQ:  $M = 22.13$ ,  $SD = 5.74$ ). A third group of 15 adults with ASD were recruited via advertisements and the Monash Alfred Psychiatry Research Centre volunteer database (4 female; age:  $M = 29.27$ ,  $SD = 9.17$  years; AQ:  $M = 28.60$ ,  $SD = 10.47$ ).

Diagnoses were of either autistic disorder (high-functioning) or Asperger's disorder. All diagnoses were according to DSM-IV-TR criteria (American Psychiatric Association, 2000) and established by a qualified clinician external to the study (psychiatrist, paediatrician or clinical psychologist). All participants were right-handed. Further demographic and clinical characteristics are reported in *Table S1*, Supplementary Material. Approval for this research was granted by the

Monash University Human Research Ethics Committee and the Alfred Hospital Ethics Committee. All participants provided informed consent.

## 2.2 Procedure

Participants sat at a desk with their right arm resting in a fixed position. A prosthetic right arm was posed in an anatomically plausible position in front of the participant, while the participant's corresponding limb was hidden from view. The limbs were spaced 20 cm apart in the horizontal plane (as measured from the middle fingers), aligned in the vertical plane, and positioned with approximately equivalent hand configuration and orientation. The prosthesis was visually similar to a human limb with respect to physical proportions, skin detail, and compression to touch. An experimenter applied stroking concurrently to each limb using a pair of soft brushes (2–2.5 x 0.5 cm tip size). Stimulation was applied to the dorsal surface of the fingers and hand. Each trial consisted of either *synchronous* or *asynchronous* stimulation, applied for 3 mins at approximately 1–2 Hz. In the synchronous stimulation condition, stroking was applied in temporal synchrony to corresponding locations of each limb. In asynchronous stimulation trials, stroking was both temporally and spatially asynchronous. Asynchronous stimulation constitutes the standard control condition in research employing the RHI, and tends not to elicit behavioural and physiological responses characteristic of the RHI (for review, see Ehrsson, 2012). Participants were instructed to attend to the stroking of the prosthetic limb during stimulation. Sixteen trials were conducted in total (eight with synchronous stimulation; eight with asynchronous stimulation). Trial order was randomised for each participant.

Each trial included pre- and post-stimulation estimates of limb position and a post-stimulation reach-to-grasp movement. At the end of each trial, participants completed a questionnaire to assess their subjective experience of the illusion. Throughout the experiment, the real and prosthetic limbs were situated in separate compartments of an observation box that spanned the length of the desk. Compartmentalised lighting allowed the prosthetic limb to be visible only in the stimulation phase of each trial, while the participant's limb was occluded

throughout the experiment. The reach target was presented only in the reaching phase of each trial. A smock blocked the participant's view of how the real and prosthetic limbs entered the box.

## **2.3 Perceptual measures**

### **2.3.1 Illusion statements**

Participants reported on their experience of the stimulation period in each trial using a questionnaire comprised of 11 statements (see Palmer et al., 2013, for full description). Three statements were worded to capture the typical phenomenological qualities of the illusion: (i) "It seemed as if I was feeling the touch of the paintbrush in the location where I saw the rubber hand being touched," (ii) "It seemed as though the touch I felt was caused by the paintbrush I could see touching the rubber hand," and (iii) "It felt as if the rubber hand was my hand." Also included were 8 control statements that were not expected to differ systematically between synchronous and asynchronous conditions (e.g., "It seemed as if I might have more than one right hand or arm"). Participants rated their agreement with each statement on a 20 cm horizontal visual analogue scale. Average ratings across illusion-related and control items were analysed. Statement order was randomised for each trial.

### **2.3.2 Illusion onset latency**

Participants pressed a footswitch during the stimulation phase of each trial when they first agreed with the statement, "It seemed as though the touch I felt was caused by the paintbrush I could see touching the rubber hand", or the statement, "It seemed as if I was feeling the touch of the paintbrush in the location where I saw the rubber hand being touched." Participants were not required to press the footswitch if they did not agree with either statement (these trials were treated as missing data; missing data due to equipment issues described below). The illusion onset latency for each trial was recorded as the duration between when the

experimenter first began applying stimulation and when the participant first pressed the footswitch.

### 2.3.3 Proprioceptive drift

Directly before and after each stimulation period, participants estimated the position of their right hand. A bar was positioned across the box, above the participant's right hand, in the horizontal plane from the participant's perspective. The experimenter slid a marker across the bar, and the participant verbally indicated when the marker was estimated to be directly above the centre knuckle of their hand. The experimenter recorded the position of the marker to the nearest millimetre. Proprioceptive drift was calculated as the difference between pre- and post-stimulation estimates, with positive values indicating drift towards the prosthetic limb.

## 2.4 Kinematic measures

Participants performed a reach-to-grasp movement in each trial using the hand that had received stimulation. The target of the movement was a 4.5 cm by 18 cm bright yellow cylinder, located 13 cm forward and 5 cm to the right of the participant's hand. Participants began the movement when a light was switched on to allow vision of the upper ~2 cm of the target; vision of their hand and the prosthetic limb was blocked throughout the movement.

Displacement was recorded continuously by an electromagnetic tracking device together with a sensor attached centrally to the dorsal surface of the hand (Ascension Technology Corporation 3DGuidance trakStar with mid-range transmitter; reported resolution of 1.4 mm and 0.5°). Recording was at 60 Hz with a 50 Hz notch filter. The three spatial dimensions of movement were the subject of analyses. All trials were visually screened for recording artefacts. Two participants were excluded from reach analyses due to extensive recording failures in these sessions; across the remaining 43 participants (688 trials), 11 trials in total were similarly excluded due to equipment issues. For each trial, movement onset and

offset were defined as when velocity first exceeded 20 mm/s for 0.05 s when proceeding anterograde and retrograde through the time series, respectively (this follows Kammers, de Vignemont, et al., 2009; Kammers, Verhagen, et al., 2009; Palmer et al., 2013). A Savitzky-Golay filter was used to smooth and then differentiate displacement data (frame length = 11–19; polynomial order = 2–3).

Area under the curve of the Euclidean jerk profile was a measure of interest. Jerk is the change in acceleration over time (the third derivative of displacement). Minimisation of mean or integrated squared jerk is a theoretical criterion for producing smooth, naturalistic point-to-point trajectories (Hogan & Flash, 1987). Previous clinical studies have used jerk measures to quantify movement performance (e.g., Nobile et al., 2011; Romero et al., 2003; Teulings et al., 1997), and we have previously demonstrated that individuals low in autistic features show increased integrated jerk in reaching movements following the RHI (Palmer et al., 2013).

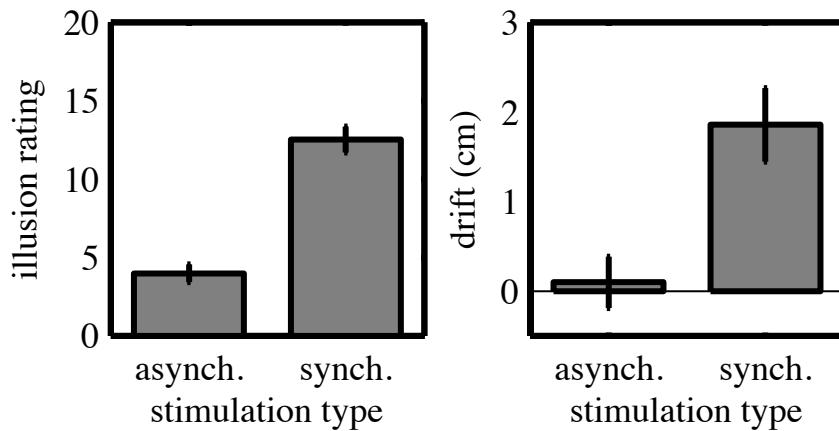
Movement sub-components were also examined to probe for features that may underlie differences in overall execution. This analysis drew on a fitting method recently developed by Friedman and colleagues to decompose recorded trajectories into constituent sub-movements (Friedman, 2012; Friedman, Brown, & Finkbeiner, 2013). This procedure finds the minimum number of (potentially overlapping) sub-movements that sum together to reproduce the observed velocity profile. Individual sub-movements are assumed to minimise jerk (i.e., show a Gaussian velocity profile) and fit certain temporal and spatial constraints. The fitting procedure described in Friedman et al. (2013) was run for 1:10 sub-movements with an error threshold of 0.03. Spatial bounds were set for x (-200–300) and y (-100–200) dimensions based on the physical proportions of the task environment. This procedure returned 2–3 sub-movements for 84% of trials, with a single trial returning no valid solution after not converging. The average reconstruction error was 0.0166 ( $SD = 0.0026$ ), similar to (Friedman et al., 2013). We examined parameters of the first two velocity sub-movements as >99% of trials contained at least two sub-movements. These parameters were peak velocity, time to peak velocity, full width at half maximum (a measure of duration)

and onset time. Fitted sub-components for a single trial are illustrated in *Figure 3a* in Results.

### 3. Results

#### 3.1 Perceptual measures

The typical perceptual effects of the RHI were observed across groups. A  $2 \times 2 \times 3$  mixed ANOVA was performed for self-reported ratings of the illusion, with statement type (control versus illusion-related), stimulation type (synchronous versus asynchronous) and group (Low AQ versus High AQ versus Clinical) as factors. Importantly, a significant interaction was observed between statement type and stimulation type,  $F(1, 42) = 113.09, p < .0001$ . Post-hoc *t*-tests indicated that illusion items were rated higher following synchronous stimulation ( $M = 12.52, SD = 5.55$ ) compared to asynchronous stimulation ( $M = 3.99, SD = 3.79$ ),  $t(44) = 11.09, p < .0001$ , Hedges'  $g_{av} = 1.81$  (*Figure 1*). Similarly, illusion items were rated higher than control items for synchronous stimulation,  $t(44) = 9.03, p < .0001$ , Hedges'  $g_{av} = 1.32$  (illusion ratings:  $M = 12.52, SD = 5.55$ ; control ratings:  $M = 6.74, SD = 3.16$ ), but not for asynchronous stimulation, for which a lesser difference in the opposite direction was observed,  $t(44) = -4.59, p < .0001$ , Hedges'  $g_{av} = 0.37$  (illusion ratings:  $M = 3.99, SD = 3.79$ ; control ratings:  $M = 5.25, SD = 3.00$ ). Together, these results indicate that the phenomenological features of the illusion typically reported in the literature tended to be experienced following synchronous stimulation but not asynchronous stimulation, as expected. A significant interaction effect was also observed between statement type, stimulation type and group,  $F(2, 42) = 4.33, p = .02$ . Post-hoc tests indicated that all three groups demonstrated the same pattern of effects as reported above for the whole sample, however (these post-hoc tests and further main effects are reported in Supplementary Material).



**Figure 1.** Perceptual measures of the RHI. Mean illusion ratings and drift in perceived arm position are shown across RHI conditions. synch., synchronous; asynch., asynchronous. Error bars indicate standard error.

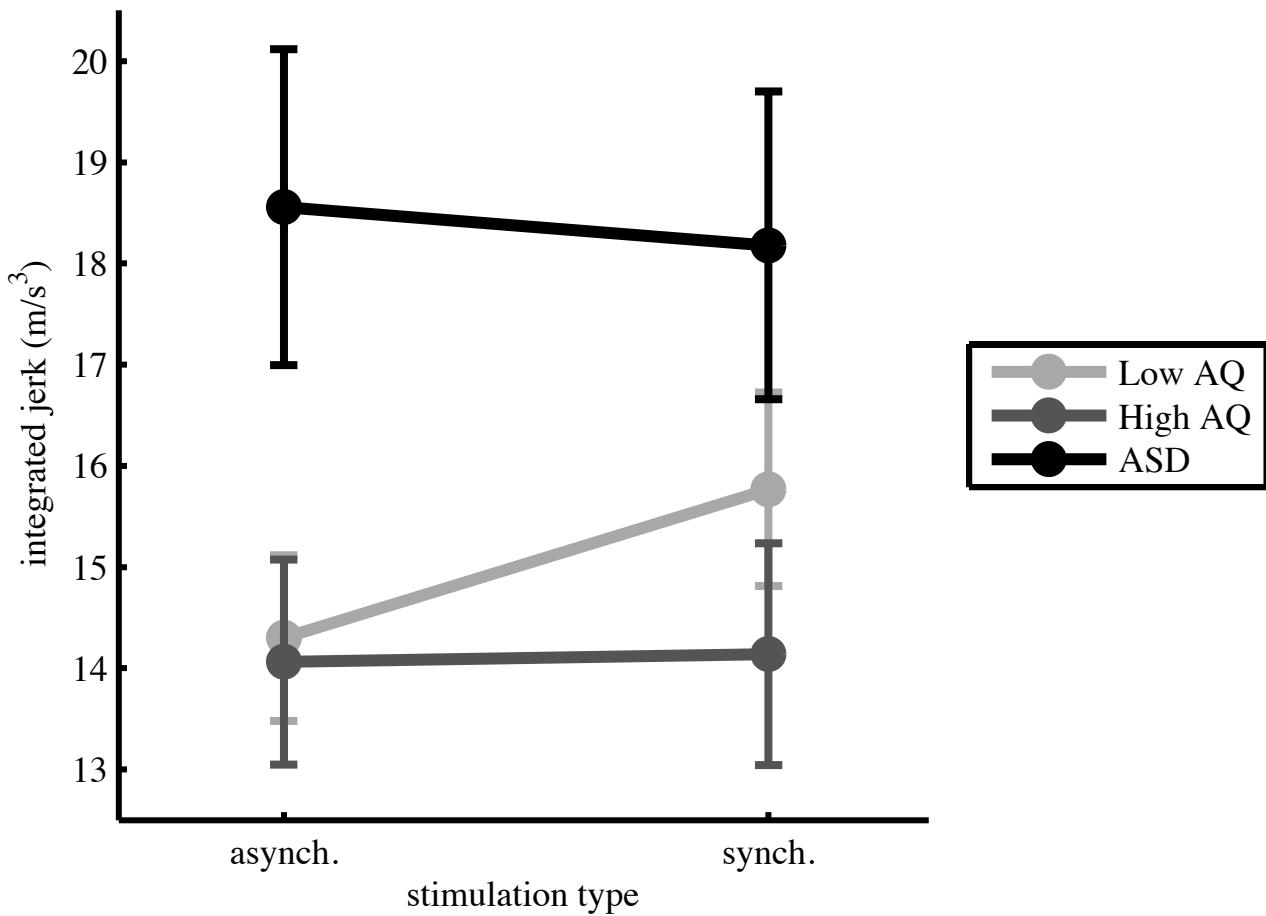
Groups similarly did not differ in the average time required to induce the illusion in the synchronous stimulation condition. A Kruskal-Wallis test with group as the factor was performed on illusion onset latency during synchronous stimulation. (A non-parametric test was used because the distribution of this data was positively skewed). There was no significant difference across the three participant groups,  $\chi^2 (2, n = 33) = 0.14, p = .93$ . Mean onset latency for the full sample was 51.25 sec ( $SD = 42.29$ ; median = 48.65; median absolute difference = 36.60).

Synchronous stimulation induced drift in perceived arm position towards the prosthetic limb, and, consistent with other perceptual measures, this effect did not differ in relation to autistic characteristics. A 2 x 3 mixed ANOVA was performed for proprioceptive drift measurements, with stimulation type (synchronous versus asynchronous) and group (Low AQ versus High AQ versus Clinical) as factors. Drift in perceived arm position towards the prosthetic limb was significantly greater for synchronous ( $M = 1.86, SD = 2.74$ ) than asynchronous stimulation ( $M = 0.10, SD = 1.91$ ),  $F(1, 42) = 43.73, p < .0001, \eta^2_p = .51, \eta^2_G = .13$  (Figure 1). No other main or interaction effects were significant ( $p > .05$ ).

### 3.2 Kinematic measures

#### 3.2.1 Integrated jerk

While the perceptual effects of the RHI were intact across groups, autistic characteristics were found to modulate reaching movements performed subsequent to the experience of the illusion. A 2 x 3 mixed ANOVA was performed for the integrated jerk index of movement performance, with stimulation type (synchronous versus asynchronous) and group (Low AQ versus High AQ versus Clinical) as factors. A significant interaction was observed between stimulation type and group,  $F(2, 40) = 5.26, p = .009, \eta^2_p = .21$  (*Figure 2*). In replication of our previous study in nonclinical individuals (Palmer et al., 2013), the Low AQ group showed increased integrated jerk in movements performed subsequent to synchronous stimulation ( $M = 15.77, SD = 3.58$ ) compared to asynchronous stimulation ( $M = 14.30, SD = 3.06$ ),  $t(13) = 3.55, p = .004$ , Hedges'  $g_{av} = 0.43$ . In contrast, the High AQ group showed no difference between synchronous ( $M = 14.14, SD = 4.10$ ) and asynchronous ( $M = 14.06, SD = 3.79$ ) stimulation conditions,  $t(13) = 0.23, p = .83$ , Hedges'  $g_{av} = 0.02$ . Moreover, in the present study we were able to extend this analysis to individuals with ASD, who similarly showed no change in integrated jerk across synchronous ( $M = 18.18, SD = 5.90$ ) and asynchronous ( $M = 18.56, SD = 6.05$ ) conditions,  $t(14) = -0.79, p = .44$ , Hedges'  $g_{av} = 0.06$ . Additionally, there was a main effect of group,  $F(2, 40) = 3.59, p = .037, \eta^2_p = .15$ . Pairwise comparisons with Bonferroni adjustment indicated that the Clinical group ( $M = 18.37, SD = 5.90$ ) exhibited significantly greater integrated jerk in their reaching movements (across conditions) than the High AQ group ( $M = 14.10, SD = 3.89; p = .015$ ). The Low AQ group did not differ significantly from either other group ( $p > .05$ ). No other main or interaction effects were significant ( $p > .05$ ).



**Figure 2.** Mean integrated jerk of reaching movements performed subsequent to the RHI. synch., synchronous; asynch., asynchronous. Error bars indicate standard error.

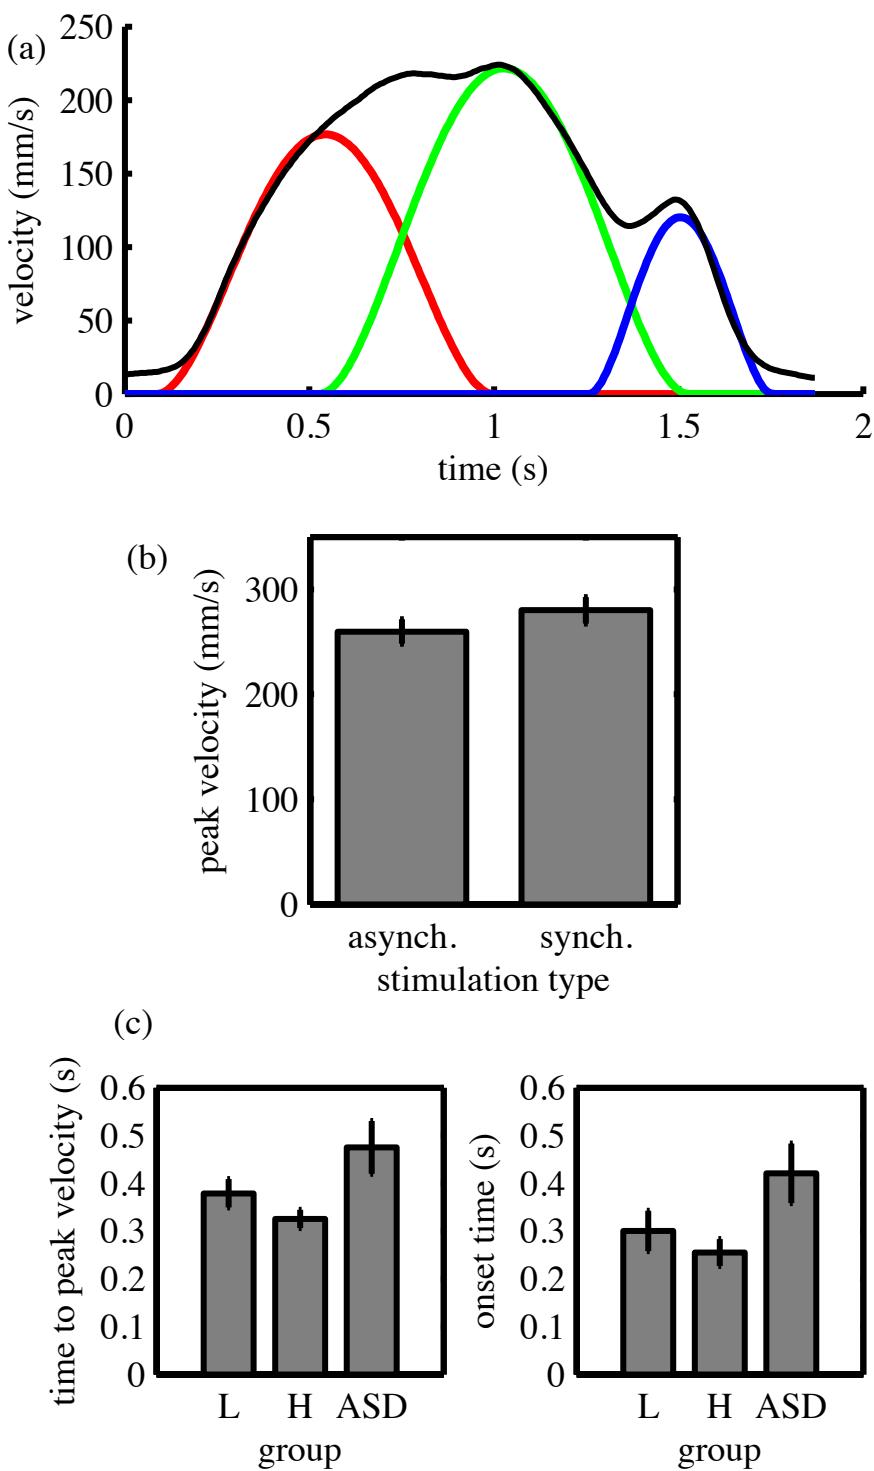
### 3.2.2 Sub-movement analysis

Modelling the constituent sub-components of the reach-to-grasp movements performed subsequent to the illusion shed light on the features of movement that likely contributed to the differences in performance noted in the previous section. Parameters of the first and second sub-movements were each analysed in a 2 x 3 mixed ANOVA with stimulation type (synchronous versus asynchronous) and group (Low AQ versus High AQ versus Clinical) as the factors.

Peak velocity of the first sub-movement did not differ across groups or RHI conditions, nor was there an interaction effect ( $p > .05$ ). However, for the *second* sub-movement, there was a significant interaction between stimulation type and group,  $F(2, 40) = 4.88, p = .013$ . There was no main effect of either factor ( $p > .05$ ). Post-hoc *t*-tests indicated that the Low AQ group showed greater peak velocity in the second sub-movement following synchronous stimulation ( $M = 280.06, SD = 46.87$ ) compared to asynchronous stimulation ( $M = 259.82, SD = 43.09$ ),  $t(13) = 3.527, p = .004$ , Hedges'  $g_{av} = 0.44$  (*Figure 3b*). In contrast, the High AQ and Clinical groups showed no difference in peak velocity between the synchronous (High AQ:  $M = 260.46, SD = 44.50$ ; Clinical:  $M = 288.96, SD = 112.95$ ) and asynchronous (High AQ:  $M = 273.35, SD = 45.50$ ; Clinical:  $M = 301.48, SD = 121.01$ ) stimulation conditions [High AQ:  $t(13) = -1.39, p = .19$ , Hedges'  $g_{av} = 0.28$ ; Clinical:  $t(14) = -1.28, p = .22$ , Hedges'  $g_{av} = 0.10$ ].

There was a significant main effect of group for time to peak velocity of the first sub-movement,  $F(2, 40) = 3.85, p = .03, \eta^2_p = .16$  (*Figure 3c*). Pairwise comparisons with Bonferroni adjustment indicated that the Clinical group ( $M = .48, SD = .21$ ) took longer to reach peak velocity of the first sub-movement than the High AQ group ( $M = .33, SD = .07; p = .009$ ). The Low AQ group did not differ significantly from either other group ( $p > .05$ ). The other main and interaction effects for this variable were non-significant ( $p > .05$ ). Time to peak velocity of the second sub-movement showed no main or interaction effects ( $p > .05$ ).

A significant main effect of group also existed for the onset time of the second sub-movement,  $F(2, 40) = 3.39, p = .044, \eta^2_p = .15$  (*Figure 3c*). Similar to the difference in time to peak velocity of the first sub-movement, pairwise comparisons with Bonferroni adjustment indicated that the onset time of the second sub-movement was later for the Clinical group ( $M = .42, SD = .24$ ) than for the High AQ group ( $M = .25, SD = .10; p = .016$ ). The Low AQ group did not differ significantly from either other group ( $p > .05$ ), and the further main and interaction effects for this variable were non-significant ( $p > .05$ ). Further analysis of the reach data is reported in Supplementary Material.



**Figure 3.** Sub-component analysis of reach-to-grasp movements. (a) Modelled sub-components of a single reach-to-grasp movement. Black, recorded data; red, first sub-movement; green, second sub-movement; blue, third sub-movement. (b) Mean peak velocity of the second sub-movement for the Low AQ group across RHI

conditions. (c) Mean time to peak velocity (first sub-movement) and onset time (second sub-movement) across participant groups. L = Low AQ; H = High AQ. Error bars indicate standard error.

#### **4. Discussion**

This study was designed to examine sensitivity to the preceding sensory context in motor function with respect to clinical and subclinical autistic features. In nonclinical individuals, those higher in autistic characteristics were less sensitive to the presence of the RHI in their reaching movements than those lower in autistic characteristics. Adults with a diagnosis of ASD similarly exhibited little difference in kinematic parameters across illusory and control conditions. Sensitivity to the illusion manifested in the Low AQ group as increased integrated jerk (in replication of our previous study of subclinical autistic characteristics; Palmer et al., 2013) and increased peak velocity in the second sub-component of movement. The latter feature is likely to contribute to or underlie increased integrated jerk, as integrated jerk is responsive to changes in the shape of the velocity profile. All groups reported the typical subjective effects of the illusion and exhibited the same degree of drift in perceived arm position towards the prosthetic limb. Thus, the observed differences in reaching movements are explained better in terms of an association between autistic features and sensitivity to prior (contextual) information rather than a general resistance to the illusion in the High AQ and Clinical groups.

This pattern of results supports and extends to movement the recent notion that the autism spectrum is characterised by reduced top-down modulation of sensory processing (Hohwy, 2013; Lawson et al., 2014; Palmer et al., 2013; Pellicano & Burr, 2012; Van de Cruys et al., 2014). In a predictive processing view of the RHI paradigm, initial expectations for arm position influenced by the illusion are updated iteratively with sensory (proprioceptive) feedback received once movement is underway. In the Low AQ group, movement in illusory and control conditions began in the same manner (as indicated by the lack of difference

between conditions in the parameters of the first sub-movement) but differed in peak velocity of the second sub-movement. This is consistent with the hypothesis that movement performance in this group is modulated by conflict between sensory feedback and illusory expectations for arm position: as movement unfolds participants in the Low AQ group accumulate sensory evidence for the true position of the arm and in that light make in-flight corrections to the movement. Following on from this, we can understand the reduced sensitivity to the illusion in the High AQ and Clinical groups as reflecting a greater weighting of sensory feedback in determining arm position during movement, such that prior representations of the environment are relatively circumscribed in their influence.

Individuals with ASD were also distinct from nonclinical individuals high in autistic traits, the former showing greater integrated jerk in movements across conditions. The sub-movement analysis was again revealing about the features of movement that may contribute to this performance difference; specifically, the Clinical group (compared to the High AQ group) showed a later time to peak velocity of the first sub-component of movement and a later onset time of the second sub-component of movement. This sluggishness in the early stage of movement coheres with previous research investigating abnormalities in movement initiation and preparation in ASD (e.g., Enticott, Bradshaw, Iansek, Tonge, & Rinehart, 2009; Rinehart, Tonge, et al., 2006). Thus, while nonclinical individuals higher in autistic traits resemble clinical individuals in their lack of motor sensitivity to the RHI, these groups are distinct in early features of movement that manifest across contexts. A regression analysis indicated no relationship across the sample between AQ score and movement performance across conditions (reported in Supplementary Material), similarly suggesting that these features of movement are specific to clinical individuals. This highlights how examining both clinical individuals and nonclinical variation in autistic traits in the same experiment can provide a more complete characterisation of how the clinical condition presents with respect to variation that exists across the general population. Furthermore, this pattern of results points to the possibility that the processing differences that contribute to reduced sensitivity to the RHI in the High AQ and ASD groups occurs

to such an extent in the latter that difficulties in motor performance are experienced across contexts.

Motor incoordination (e.g., clumsiness) occurs commonly in ASD across a range of different motor behaviours, including reaching movements and gait (Fournier et al., 2010; Gowen & Hamilton, 2013). The results of the present study furnish an account of the motor symptoms of ASD in terms of the relative weighting of prior or contextual information against sensory feedback. A fundamental assumption of the Bayesian approach to perception is that sensory information is noisy and ambiguous, such that drawing upon prior and contextual information is necessary to determine the state of the external world. Thus, the increased weighting of sensory information in perceptual inference that is suggested to occur in ASD leads directly to an account of motor incoordination in terms of how the brain estimates the state of the body during movement. Namely, relying too highly on the incoming sensory information at the expense of prior information should typically lead to a *less* accurate sense of body position, which may contribute to clinical symptoms of motor incoordination (and the reduced smoothness of movement observed for the clinical group in the present study). Moreover, while the context of the illusion misleads performance in the RHI, sensitivity to higher-order contextual information may more commonly be of benefit to accurate motor performance (e.g., when performing movement without visual feedback, or when contextual factors like weight on the arm modulate the relationship between actions and their sensory consequences).

The concepts that we draw upon to elucidate differences in sensitivity to the RHI across the groups may also be useful in understanding the observed differences in movement initiation in ASD. In their application of predictive processing to action, Friston and colleagues have emphasised the role of *sensory attenuation* in movement initiation (Brown et al., 2013). In brief, their notion is that action comes about when predictions regarding the flow of proprioceptive input are fulfilled by peripheral responses that engage the muscles to change the bottom-up signal (Adams, Shipp, et al., 2013; Shipp et al., 2013). This contrasts with the more passive process of updating predictions to match input suggested to occur in the

perceptual system. To begin movement, sensory evidence for the hypothesis concerning the current (true) arm position must be down-weighted such that alternate hypotheses regarding arm position are favoured. As noted, the mechanism thought to underlie a reduced influence of prior or higher-level expectations in autistic perception is a tendency to weight the sensory input highly relative to top-down predictions (i.e., increased gain on prediction errors) (Lawson et al., 2014; Palmer et al., 2013; Van de Cruys et al., 2014). An implication of increased sensory weighting may be that the attenuation of sensory evidence that is suggested to facilitate movement initiation is compromised, contributing to differences in the early stages of movement.

Is increased sensory weighting in ASD contextually driven or a chronic feature of sensory processing? In the predictive processing framework, the relative weighting of sensory information against prior expectations is adjusted top-down in a state-dependent manner in response to changes in the expected precision of sensory signals (i.e., the estimated uncertainty of the environment). This is a mechanism that has the potential to add further nuance regarding how subtle differences in sensory processing may manifest in a complex manner across contexts, which may be crucial in accounting for the complex pattern of sensory differences in ASD and the heterogeneity in symptoms reported between and within individuals with ASD. The results of the present study point to both *context-independent* differences in ASD, exhibited in movement performance across conditions, and *context-dependent* differences in ASD, exhibited in differing responses to the illusion in movement. While there is evidence that proprioceptive estimates in non-illusory conditions are no more accurate or precise in ASD than in controls (Fuentes et al., 2011), the present study differs in incorporating an uncertain context in which conflict between cues for arm position derived from the illusion and proprioceptive input during illusion induction may induce an expectation for low precision in the sensory input. That is, sensorimotor input can normally be interpreted unequivocally under long-held, very stable expectations about body-image, body-schema and bodily self-awareness but the RHI challenges these expectations and throws doubt on the sensory input. Those higher in the autism spectrum may therefore weight sensory input more strongly than others in

particular contexts due to a reduced response to cues that suggest that sensory input should be distrusted.

Importantly, we can begin to see how differences in the depth of the cortical hierarchy through which updating occurs could underlie clinical and subclinical autistic features. Specifically, we can situate nonclinical individuals lower in autistic traits as tending to appeal more to higher-level causes to explain sensory input when compared to those higher in autistic features, while clinical individuals have a tendency to predominantly recruit levels that are lower again. It is clear that differences in this regard could be adaptive (as in nonclinical variation in autistic features) or not (as in the clinically-defined condition), given the equivocal challenge of determining where in the causal hierarchy to account for changes in sensory input from within the skull. We can speculate that differences in the recruitment of higher levels in the hierarchy may be reflected in neurobiological features found in ASD such as reduced long-range connectivity (e.g., Just, Cherkassky, Keller, & Minshew, 2004) and greater intra-individual variability in evoked cortical responses (Dinstein et al., 2012). Since representations of context help to explain fluctuations in sensory input over relatively longer time scales, more circumscribed context sensitivity in autistic sensory processing could contribute more generally to reduced social comprehension, sensory impairments and a stronger desire for predictability and routine.

There is evidence that the perceptual experience of the RHI is facilitated by prior expectations regarding bodily representation (e.g., Hohwy & Paton, 2010; Tsakiris & Haggard, 2005). That the typical perceptual effects of the RHI were exhibited by the ASD group thus suggests that individuals with ASD are able to learn informative priors but differ instead in the relative weighting of priors against conflicting sensory signals (for discussion, see Brock, 2012; Lawson et al., 2014; Skewes et al., 2014). This is consistent with the mixed evidence on visual illusions in ASD, which tends to suggest that prior information influences visual perception in ASD but to a lesser extent than controls (Mitchell et al., 2010). The differences in motor behaviour following the illusion despite the similar perceptual experience across groups points to how subtle atypicalities in the integration of prior

expectations with sensory information may manifest differently across tasks depending on factors such as the nature of the priors involved (e.g., expectations regarding bodily representation developed over long time scales versus shorter-term contextual information regarding body position during movement). Similarly, the predictive processing account emphasises the continual and dynamic integration of incoming sensory signals with existing expectations regarding the causes of sensory input, which the reaching task used in the present experiment is likely to be more sensitive to than perceptual measures of the illusion due to the immediate demand of integrating sensory feedback with existing estimates of arm position as the movement unfolds. Thus, it may be important for future research in this area to similarly consider the temporal nature of perceptual inference and include measures that are sensitive to this process.

Reduced sensitivity to the RHI in movement fits broadly with established theories of autistic perception that suggest reduced sensitivity to more global or contextual information (weak central coherence, WCC; Happé & Frith, 2006) or enhanced lower-level perceptual functioning (EPF; Mottron et al., 2006). Specifically, WCC may predict a reduced sensitivity to the illusion in general due to reduced global integration of sensory information, while EPF may predict better movement performance in the context of the RHI due to enhanced discriminability of proprioceptive input (though previous research has not supported better proprioceptive discrimination in ASD, Fuentes et al., 2011). The predictive processing account has an advantage in this regard in accounting for why the RHI is experienced similarly in the ASD group but modulates subsequent movement differently to low AQ controls (as discussed earlier), and further, why movement performance is generally poor in the ASD group in addition to this group showing insensitivity to the context of the RHI. Inferential accounts of autistic perception formalise the distinction between bottom-up and top-down sensory processes (Pellicano & Burr, 2012) within a computational framework, and situate clinical symptoms and related nonclinical variation within a general model of brain function. Drawing on the biological and computational depth of inferential models may be important for elucidating the mechanisms underlying autistic symptoms; for example, the concept of weak central coherence can be cast in predictive

processing terms, leading to corresponding hypotheses regarding the cortical circuitry and neurotransmitter systems involved (Lawson et al., 2014; Van de Cruys et al., 2014). Further research may also be able to distinguish predictive processing accounts from other theories by examining the role of sensory uncertainty in modulating differences in perceptual or sensorimotor outcomes between ASD and controls.

In this study, we find that both individuals with ASD and nonclinical individuals high in autistic traits show reduced sensitivity to a multisensory illusion of limb ownership in subsequent reaching movements despite experiencing the perceptual effects of this illusion. In addition, clinical and nonclinical individuals high in autistic traits are distinguishable in the integrated jerk and earlier phases of movement across conditions. These sensorimotor differences can be understood in terms of the relative weighting of sensory feedback against existing expectations for body position as movement unfolds; these results thus extend recent predictive processing accounts of ASD to sensorimotor function and to differences across the nonclinical population.

## Acknowledgements

The authors wish to thank Uta Frith for helpful comments on this manuscript, and Owen Hammond and colleagues at the Monash Instrumentation Facility for creating the mirror box apparatus used in this experiment. This work was funded by an Australian Research Council Discovery Grant (DP1311336). J.H. is supported by an Australian Research Council Future Fellowship (FT100100322). P.E. is supported by a NHMRC Career Development Fellowship (GNT1052073). B.P. is in part supported by a Monash Inter-Disciplinary Research Grant, 2013.

## S1. Supplementary methods

**Table S2.** Group demographics

	Low AQ	High AQ	Clinical
<i>n</i>	15	15	15
Age (years)	30.20 (7.31)	29.87 (8.61)	29.27 (9.17)
Gender (f:m)	8:7	5:10	4:11
AQ			
Sociability	19.27 (3.52)	25.73 (6.89)	28.47 (8.55)
Mentalising	8.47 (1.64)	11.60 (3.31)	14.87 (4.32)
Detail Orientation	10.53 (3.02)	16.67 (3.92)	17.60 (3.85)
Total (likert-scored)	86.80 (10.08)	117.60 (11.64)	132.87 (24.68)
Total (binary-scored)	8.07 (3.96)	22.13 (5.74)	28.60 (10.47)
RAADS-R			
Social Relatedness	-	-	23.93 (12.21)
Circumscribed	-	-	45.73 (16.72)
Interests			
Sensory Motor	-	-	24.40 (10.74)
Social Anxiety	-	-	20.67 (10.17)
Total	-	-	113.40 (34.30)
SRS-2			
Social Awareness	-	-	59.27 (12.09)
Social Cognition	-	-	62.40 (12.97)
Social	-	-	61.80 (16.32)
Communication			
Social Motivation	-	-	62.60 (13.36)
RIRB	-	-	67.67 (13.89)
Total	-	-	64.93 (13.29)

Means and standard deviations are shown for continuous variables. Abbreviations: AQ: Autism-Spectrum Quotient (Baron-Cohen et al., 2001) (see Palmer, Paton, Enticott, and Hohwy (2015) for description of the subscales reported here); RAADS-R: Ritvo Autism Asperger Diagnostic Scale-Revised (Ritvo et al., 2011); RIRB: Restricted Interests and Repetitive Behavior subscale; SRS-2: Social Responsiveness Scale, Second Edition (Constantino & Gruber, 2012) (SRS T-scores from self ratings are reported here). Four individuals in the clinical group reported co-morbidities (1 depression; 1 anxiety disorder; 1 depression/anxiety disorder; 1 depression/anxiety disorder/dissociative disorder/dyslexia). Co-morbid depression and anxiety are common in ASD (Matson & Williams, 2014), thus the sample is likely better representative of the general ASD population with these individuals included. Four clinical participants were medicated (1 serotonin-norepinephrine reuptake inhibitor; 1 lithium/atypical antipsychotic/benzodiazepine; 1 selective serotonin reuptake inhibitor/typical antipsychotic/benzodiazepine; 1 atypical antipsychotic/anticholinergic).

## S2. Supplementary results

### S2.1 Illusion ratings

Further to the three-way interaction between statement type, stimulation type and group reported for illusion ratings in the main text: Post-hoc t-tests indicated that all three groups showed (i) increased illusion ratings following synchronous stimulation compared to asynchronous stimulation, and (ii) that illusion ratings were greater than control ratings for synchronous stimulation but not for asynchronous stimulation. These tests are reported below in *Table S2*.

In addition to the interaction effects for illusion ratings reported in the main text, there was a main effect of statement type,  $F(1, 42) = 45.73, p < .0001, \eta^2_p = .52$ . Illusion statements ( $M = 8.26, SD = 4.00$ ) were rated higher than control statements ( $M = 6.00, SD = 2.95$ ). There was also a main effect of stimulation type,  $F(1, 42) = 127.64, p < .0001, \eta^2_p = .75$ . Synchronous stimulation ( $M = 9.63, SD = 3.98$ ) was associated with higher statement ratings than asynchronous stimulation ( $M = 4.62, SD = 3.30$ ). All other main and interaction effects were non-significant ( $p > .05$ ).

### S2.2 Further reach analysis

The full width at half maximum parameter was not associated with significant main or interaction effects for either sub-movement ( $p > .05$ ).

To further examine the relationship between autistic characteristics and the features of movement that differed between the clinical and nonclinical (High AQ) groups across the illusion and control conditions, a standard regression was performed with AQ score as the dependent variable and movement features (integrated jerk, time to peak velocity of the first sub-movement, and onset time of the second sub-movement) as the independent variables. Unsurprisingly, given their mutual relationship to movement initiation, there was strong evidence for multicollinearity between the time to peak velocity and onset time variables based

on their Pearson's correlation ( $>.7$ ), tolerance ( $<.10$ ) and variance inflation factor ( $>10$ ). The onset time variable was thus omitted from the regression analysis. The scatterplot of standardized residuals and Mahalanobis distances indicated a single extreme outlier, which was removed from the analysis. The total variance explained by the model was small and non-significant,  $R^2 = 9.3\%$ ,  $F(2, 39) = 2.01$ ,  $p = .15$ . Neither integrated jerk ( $\beta = .20$ ,  $p = .19$ ) nor time to peak velocity of the first sub-movement ( $\beta = .22$ ,  $p = .15$ ) made a significant unique contribution. Thus, while these features of movement differed across groups as described earlier, there was not evidence that they shared a linear relationship with AQ scores across the sample when combining clinical and nonclinical individuals.

Peak lateral displacement (mm) away from the prosthetic arm was examined with a  $2 \times 3$  mixed ANOVA containing stimulation type (synchronous versus asynchronous) and group (Low AQ versus High AQ versus Clinical) as factors. Synchronous stimulation was associated with greater peak lateral displacement ( $M = 74.90$ ,  $SD = 14.58$ ) than asynchronous stimulation ( $M = 71.00$ ,  $SD = 15.05$ ),  $F(1, 40) = 9.56$ ,  $p = .004$ ,  $\eta^2_p = .19$ ,  $\eta^2_G = .02$ . No other main or interaction effects were significant ( $p > .05$ ).

*Table S3.* Post-hoc *t*-tests for three-way interaction

	<i>M</i> ( <i>SD</i> )	<i>t</i> (df)	<i>p</i>	Hedges' <i>g</i> <sub>av</sub>
<i>Low AQ</i>				
Illusion ratings (Synch)	13.81 (5.04)			
Illusion ratings (Asynch)	3.28 (3.20)	9.20 (14)	<.0001	2.49
Illusion ratings (Synch)	13.81 (5.04)			
Control ratings (Synch)	6.44 (2.50)	7.66 (14)	<.0001	1.90
Illusion ratings (Asynch)	3.28 (3.20)			
Control ratings (Asynch)	4.82 (2.26)	-3.87 (14)	<.01	0.55
<i>High AQ</i>				
Illusion ratings (Synch)	9.82 (6.10)			
Illusion ratings (Asynch)	3.74 (4.24)	4.94 (14)	<.001	1.14
Illusion ratings (Synch)	9.82 (6.10)			
Control ratings (Synch)	6.29 (3.70)	3.97 (14)	<.01	0.70
Illusion ratings (Asynch)	3.74 (4.24)			
Control ratings (Asynch)	4.56 (3.46)	-1.69 (14)	.11 (n.s.)	0.21
<i>Clinical</i>				
Illusion ratings (Synch)	13.93 (4.75)			
Illusion ratings (Asynch)	4.95 (3.93)	6.32 (14)	<.0001	2.01
Illusion ratings (Synch)	13.93 (4.75)			
Control ratings (Synch)	7.50 (3.26)	5.13 (14)	<.001	1.56
Illusion ratings (Asynch)	4.95 (3.93)			
Control ratings (Asynch)	6.38 (3.03)	-2.61 (14)	<.05	0.40

## Linking text between chapters 3 and 4

In *Chapters 2 and 3*, we reported that the experience of the multisensory rubber hand illusion differed in individuals with ASD and in relation to nonclinical variation in autism-like characteristics. These differences were reflected primarily in reaching movements performed subsequent to the illusion, with autistic features associated with reduced sensitivity to the context of the illusion in movement. These findings can be understood as reflecting the level of integration of prior or contextual expectations regarding arm position with incoming proprioceptive signals. Reduced context sensitivity in movement is thus consistent with an overweighting of sensory signals during Bayesian inference that is hypothesised to occur in ASD (Lawson et al., 2014; Palmer et al., 2013; Pellicano & Burr, 2012; Van de Cruys et al., 2014).

An attraction of drawing on Bayesian probability theory to characterise differences in perceptual or cognitive function in neuropsychiatric conditions is the promise of quantifying such differences in more precise computational terms. For instance, Bayesian modelling techniques have provided evidence that the human brain integrates crossmodal sources of information (e.g., vision and audition, or vision and proprioception) in a manner sensitive to the relative reliabilities of each unimodal source of information (Alais & Burr, 2004; van Beers, Sittig, & Gon, 1999). Furthermore, implicit prior expectations that the sensory system brings to bear during perception can also be modelled – for instance, the assumed direction of lighting sources during object shape perception. This method has been used to assess differences between individuals in their apparent implicit prior expectations that influence perceptual experience (Mareschal, Calder, & Clifford, 2013), and used for tracking changes in such expectations within an individual over time in response to experience (Adams et al., 2004). The rubber hand illusion is an attractive paradigm for examining perceptual inference, as it involves integration between conflicting sensory cues and prior expectations regarding the body (Hohwy, 2013; Hohwy & Paton, 2010; Limanowski & Blankenburg, 2013; Tsakiris & Haggard, 2005). Nevertheless, there is some distance between the

perceptual and kinematic measures of the illusion that we employ in *Chapters 2 and 3* and the predictive processing mechanisms that we appeal to when discussing why differences in these measures emerged across groups.

As one approach to more directly assessing the integration of incoming information with existing expectations, we examined *statistical learning*: how individuals learn about probabilistic regularities in the environment as they sample the state of the environment over time. This type of learning (or inference) on the basis of ambiguous information has been examined previously in a Bayesian framework, with this research demonstrating that the brain is able to track the uncertainty of environmental outcomes in an optimal manner (e.g., Behrens et al., 2007; McGuire, Nassar, Gold, & Kable, 2014). An advantage of this type of learning paradigm is that predictions and prediction errors can be quantified and tracked over time as new information is gathered, and moreover, the experimenter can precisely control the variance of the sampled data. In this way, it can be more directly examined how an individual uses each new piece of information to update their beliefs about the environment, and how this depends upon the noisiness of the sampled data.

In *Chapter 4*, we employ a novel statistical learning task to test the hypothesis that autistic characteristics are associated with a tendency to weight new information more highly relative to prior expectations. Finer-grained hypotheses regarding the information processing differences that characterise autistic brain function can also be tested in the domain of statistical learning: for instance, expectations about the *volatility* of environmental states can be modelled (Mathys et al., 2011; Mathys et al., 2014). More complex Bayesian modelling analyses of this nature are ongoing and not reported in the present thesis.

## Declaration for Thesis Chapter 4

### Declaration by candidate

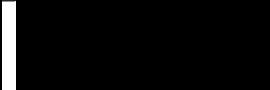
In the case of *Chapter 4*, the nature and extent of my contribution to the work was the following:

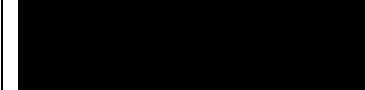
Nature of contribution	Extent of contribution (%)
<ul style="list-style-type: none"> <li>Contributed to the experiment design.</li> <li>Performed data analysis.</li> <li>Contributed to the interpretation of results.</li> <li>Wrote the paper.</li> <li>Performed data collection and analysis for pilot experiments during development of the experimental paradigm.</li> </ul>	60%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Bryan Paton	Contributed to the experiment design, interpretation of results, and writing.	N/A
Peter Enticott	Contributed to the interpretation of results and writing.	N/A
Jakob Hohwy	Contributed to the experiment design, interpretation of results, and writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date 16/02/2016
-----------------------	---	--------------------

Main Supervisor's Signature		Date 16/02/2016
-----------------------------	---	--------------------

# Chapter 4

## Developing Bayesian accounts of autism: Evidence from statistical learning

Colin J. Palmer<sup>1</sup>, Bryan Paton<sup>1,2,3</sup>, Peter G. Enticott<sup>4</sup>, Jakob Hohwy<sup>1,2</sup>

<sup>1</sup> *Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup> *ARC Centre of Excellence for Integrative Brain Function*

<sup>3</sup> *School of Psychology, Monash University, Clayton, VIC 3800, Australia*

<sup>4</sup> *Cognitive Neuroscience Unit, School of Psychology, Deakin University, 221 Burwood Highway, Burwood, VIC 3125, Australia*

## Abstract

Advances in the neurocognitive understanding of autism spectrum disorder (ASD) can be made by drawing on Bayesian theories of brain function. An emerging hypothesis in this regard is that ASD is characterised by an increased reliance on new sensory data relative to prior expectations in the information processing that underlies perception and cognition. In predictive processing models of brain function, this equates to a greater weighting of prediction error in updating predictions about the environment. The present study examines behaviour in a spatial prediction task in which prediction errors and prediction updates can be directly quantified. We report evidence that autistic features across the general population are not related to differences in the average weighting of prediction errors, nor in how this weighting is adjusted as the precision of feedback data changes. These findings are valuable in constraining the aspects of predictive processing mechanisms that are likely to be atypical in ASD. One implication is that differences in information processing mechanisms that characterise autistic perception may not extend to more explicit statistical learning; another prospect is that the gain on prediction errors in the autistic brain is characterised by a more subtle atypicality than chronic over-weighting.

## 1. Introduction

Bayesian theories of brain function provide a framework for understanding many aspects of cognition and behaviour, and very recently have been used to generate novel hypotheses regarding the neurocognitive bases of ASD (e.g., Palmer, Paton, Kirkovski, et al., 2015; Pellicano & Burr, 2012; Quattroki & Friston, 2014; Van de Cruys et al., 2014). This general approach to understanding cognition develops in part from the view of perception as a process of *implicit inference* on the environmental causes of sensory signals (Gregory, 1980; Helmholtz, 1860). In perceptual inference, the ambiguity inherent in sensory data is resolved by drawing on *prior beliefs* about the statistics of the world, such that the brain is able to represent the most likely state of the environment given the sensory data. Bayesian probability theory describes the optimal method for updating representations of the environment in light of ambiguous sensory data (Kersten et al., 2004; Vilares & Kording, 2011). Formally, prior beliefs are represented as probability distributions over hypothesised causes of sensory data, and are updated by *likelihood* functions that represent the ambiguous information carried by new sensory data. In accordance with the principles of probability theory, the degree of uncertainty associated with each of these two sources of information determines their relative influence on the updated belief.

*Predictive processing* is a neurocognitive theory of how sensory information is processed across the cerebral cortex in a manner consistent with Bayesian principles (Clark, 2013; Friston, 2005; Hohwy, 2013; Mumford, 1992). Prior beliefs furnish *predictions* about the most likely state of the environment; these predictions are compared to the actual sensory input to produce *prediction error*. Prediction error is used to update predictions in an iterative manner, thus bringing the brain's representation of the world closer to the (hidden) causes of its input. This process is made probabilistic by *weighting* prediction error according to a ratio of the precision of sensory data to the precision of existing beliefs (under Gaussian assumptions). This weighting ratio is the *learning rate*, and sets the extent to which each new data point updates the existing representation of the environment. For instance, in contexts in which sensory data are highly precise,

prediction errors are weighted highly such that fluctuations in the sensory signal more readily drive changes to the represented state of the world at the expense of prior learning. In this way, the precision-weighted updating of predictions captures core aspects of Bayesian inference. This process is suggested to occur continuously and recurrently in the form of hierarchical neuronal message passing in the cerebral cortex, with cortical feedback connections mediating predictions and feedforward connections signalling prediction error.

Atypicalities in how the weighting of prediction error is adjusted in response to context-dependent uncertainty are suggested to underlie core aspects of ASD (Hohwy, 2013; Lawson et al., 2014; Palmer, Paton, Kirkovski, et al., 2015; Pellicano & Burr, 2012; Quattroki & Friston, 2014; Van de Cruys et al., 2014) as well as other neurological and psychiatric conditions (Friston, Stephan, Montague, & Dolan, 2014). The most specific proposals at present are that autistic symptoms arise from a *chronically reduced precision of prior beliefs* in perceptual inference (Pellicano & Burr, 2012), or, similarly, a *high and inflexible weighting of prediction errors* (Van de Cruys et al., 2014). Each of these hypotheses equate to a persistently high learning rate in prediction-error updating; that is, inference that is driven to a greater extent by incoming sensory signals rather than past experience. These are exciting developments in ASD research as these theories are able to tie together diverse cognitive, psychophysical and neurophysiological characteristics of the condition, such as hypersensitivities to sensory stimulation, atypical balance between lower- and higher-level perceptual processes and reduced susceptibility to perceptual illusions (for reviews, see Lawson et al., 2014; Pellicano & Burr, 2012; Van de Cruys et al., 2014); however, empirical data that directly tests the principal hypotheses of these accounts has only just begun to emerge (Skewes et al., 2014).

In the present study, we tested the hypothesis that autistic features are associated with a persistently high learning rate in the context of statistical learning that requires inference on the underlying causes of observed data. Participants made explicit predictions regarding the location of upcoming visual markers that were sampled from (hidden) Gaussian distributions. Predictions could be adjusted trial

by trial in response to each newly revealed data point; in this way, we could directly quantify prediction errors and their influence on how predictions were updated. Individual differences in the learning rate were examined with respect to high- and low-autistic traits in non-clinical adults.

We also examined how the learning rate corresponded to the state-dependent *uncertainty* of observed data by manipulating the variance of the underlying distributions that generated this data. An important prediction of the proposal of Van de Cruys et al. (2014) is that inflexibly high weighting of prediction errors in ASD is especially problematic in contexts of greater sensory ambiguity (in which a Bayes-optimal observer would downweight prediction errors). We therefore tested the hypothesis that stronger autistic features are associated with a tendency to weight prediction errors highly regardless of the degree of sensory uncertainty, while weaker autistic features are associated with a tendency to flexibly reduce the weighting of prediction errors when sensory uncertainty increases.

## 2. Method

### 2.1 Participants

Participants were 40 adults recruited from a general university population, reporting no psychiatric or neurological diagnoses. A median split on Autism Spectrum Quotient (AQ; described in the next section) was performed to divide this sample into a *High AQ* group and a *Low AQ* group. Demographics, including AQ statistics, are reported in *Table 1*. Each participant provided informed consent and approval for this study was gained from the Monash University Human Research Ethics Committee.

**Table 4.** Sample demographics

Group	AQ score (Likert)			AQ score (binary)			Age (years)			Sex (f:m)	Handedness (r:l)
	<i>M</i>	<i>SD</i>	range	<i>M</i>	<i>SD</i>	range	<i>M</i>	<i>SD</i>	range		
High AQ	120.55	7.04	111–135	22.60	3.33	17–29	26.30	6.63	18–38	10:10	20:0
Low AQ	101.65	6.10	88–110	13.10	4.09	5–20	30.65	13.26	18–61	12:8	19:1
Full sample	111.10	11.57	88–135	17.85	6.06	5–29	28.47	10.58	18–61	22:18	39:1

Autism Spectrum Quotient (AQ) scores used in the median split and statistical analyses were those computed with Likert (4-3-2-1) rather than binary scoring (1-1-0-0), to better capture individual differences. The two AQ groups differed significantly in AQ score,  $t(38) = 8.06, p < .001, d = 2.61$ , but not age,  $t(28) = 1.31, p = .2, d = 0.43$ , nor sex,  $\chi^2 (1, n = 40) = 0.40, p = .53$ , phi = .10.

## 2.2 AQ

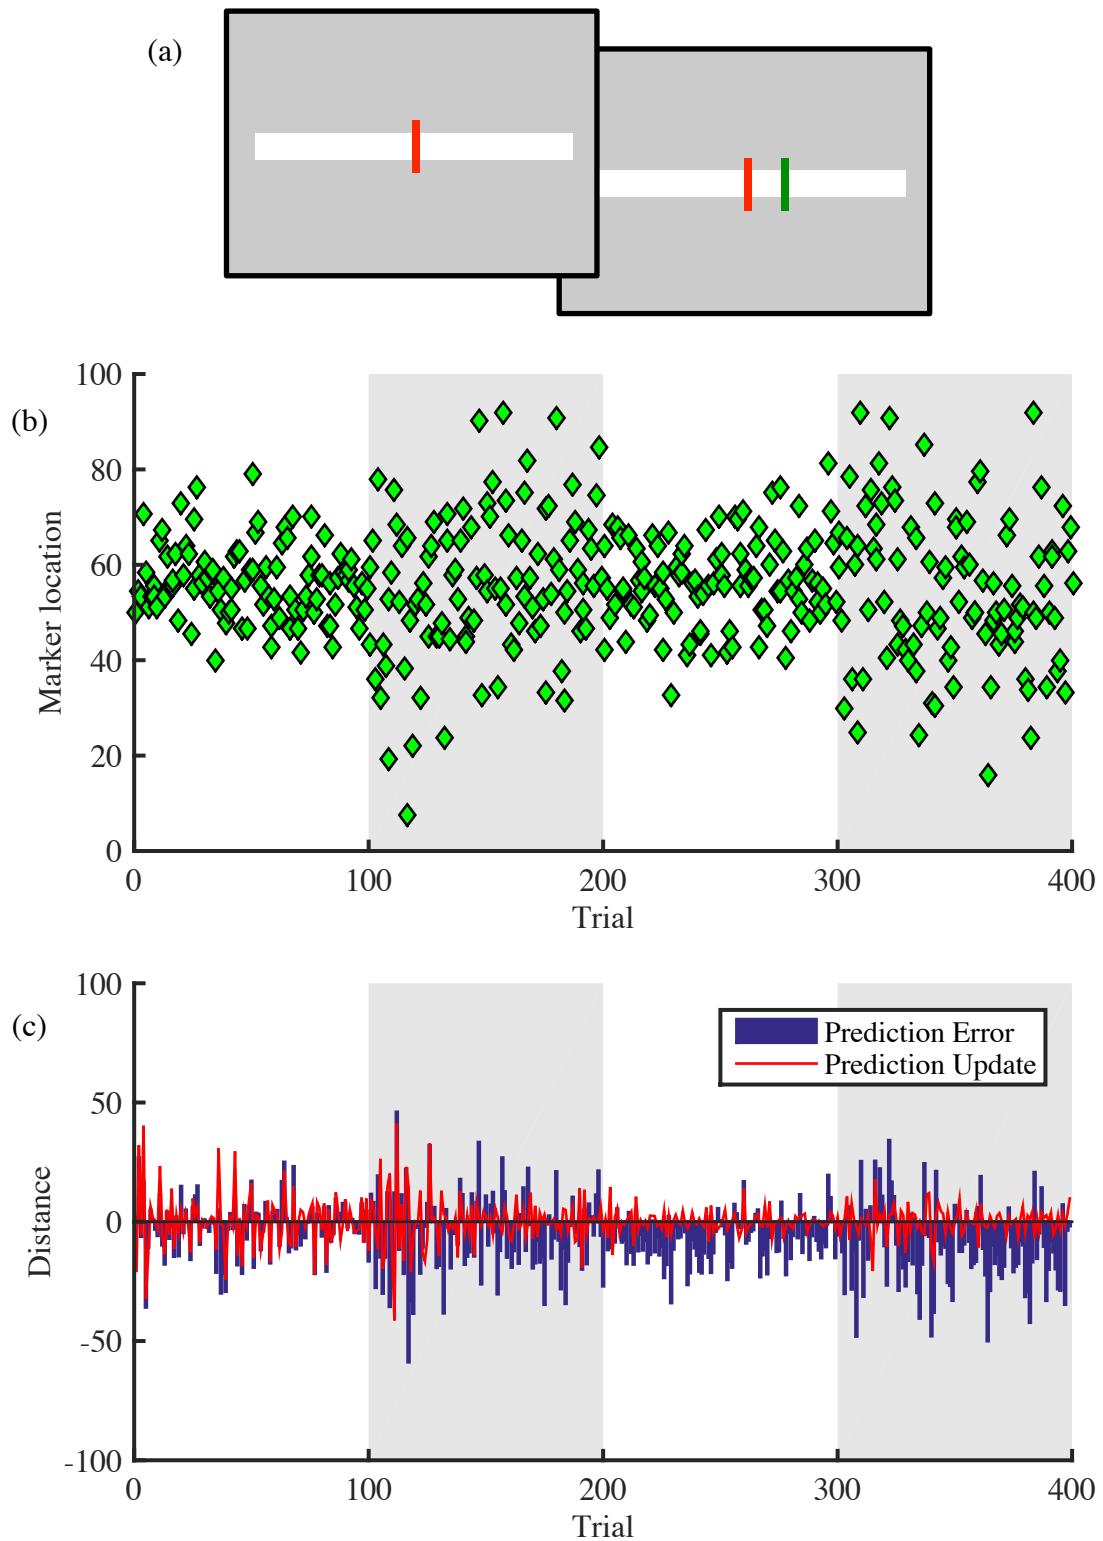
The AQ is the most well studied assessment of autistic traits in adult community samples (Baron-Cohen et al., 2001). It consists of 50 items that participants rate for personal applicability on a 4-point Likert scale; for instance, “When I’m reading a story, I find it difficult to work out the characters’ intentions.” These items are designed to assess a range of social and non-social behaviours that are characteristic of ASD, including social preferences, social cognitive abilities, restricted interests, and attentional habits. In addition to total score, the present study analysed scores from three subscales consisting of 6–11 items each: Sociability, Mentalising and Detail Orientation. These subscales were defined on an empirical basis in a recent study of the AQ in a large sample of adults drawn from the general population (Palmer, Paton, Enticott, et al., 2015). These subscales differ from the five subscales more commonly used in research, which are defined on a conceptual rather than empirical basis (Baron-Cohen et al., 2001). Higher scores for each subscale (and total score) indicate greater similarity to ASD. Adults with a diagnosis of ASD score highly on the AQ (e.g., Baron-Cohen et al., 2001; Hoekstra, Bartels, Cath, & Boomsma, 2008), and importantly, in non-clinical samples, AQ scores have been found to vary together with a range of psychophysical, cognitive and neurophysiological features characteristic of ASD (e.g., Grinter, Maybery, et al., 2009; Nummenmaa, Engell, von dem Hagen, Henson, & Calder, 2012; Walter et al., 2009).

## 2.3 Visual inference task

Participants completed a computer-based task run with Psychtoolbox v3.0.12 in MATLAB R2014b. In each trial of the task, participants were presented with a horizontal line, along which a *green target marker* would appear. The object of the task was to estimate *where* the target marker was most likely to appear in each trial. Participants were instructed that the location of the target marker would change between trials, but that there would be some consistency in where the target marker appeared across trials. The main phases of each trial are depicted in *Figure 1a*. The horizontal line was located centrally and spanned 50% of the screen

(1680 by 1050 pixel resolution). Using the mouse, participants' moved a red marker along the horizontal line and clicked to make their prediction. The starting point of the red marker was randomized in each trial. Participants then rated their confidence in their prediction being near to where the target marker would appear; ratings were made on a vertical continuous gauge (spanning 0% to 100%) that appeared after each prediction. Participants were given no time limit in making their prediction or rating their confidence. Following the confidence rating, the target marker appeared for 2 seconds. The participants' prediction (red marker) was displayed throughout the confidence rating and feedback phases such that participants could directly compare their prediction to the actual target location. A fixation cross was presented centrally between trials for 1 second, followed by a random visual noise mask for 0.5 seconds. Predictions and actual locations were recorded to the nearest 0.01% of the horizontal line length. Confidence data are not reported in the present paper.

The task consisted of four blocks of 100 trials each. In each block, target marker locations were sampled from a Gaussian distribution with a standard deviation of either one-sixth (High Variance) or one-twelfth (Low Variance) of the horizontal line length. The mean of the Gaussian distributions did not change between blocks, but was randomised across participants within the central third of the horizontal line. Target locations drawn from the Gaussian distributions were resampled if falling outside of the horizontal line. Block order was either High-Low-High-Low or Low-High-Low-High, counterbalanced across participants. The frequency of different block orders was close to even between participant groups (9:11 for Low AQ group, 11:9 for High AQ group). Target locations were sampled anew each time the task was run such that each participant received a different set of locations. An example of target locations for one participant is shown in *Figure 1b*. Participants were not instructed regarding how target markers were generated nor were they told that the method by which target markers were generated changed throughout the task. Similarly, block changes were not indicated to the participant and there were no scheduled breaks in the task.



**Figure 1.** (a) Schematic of the main phases in the computer-based inference task. Participants first predict the horizontal location at which a target marker will appear by using a mouse to position the red marker. At the end of each trial, the (green) target marker is shown. The relative size of screen components has been adjusted in these images to improve visibility. (b) An example of target marker

locations across trials for one participant. Shaded areas indicate High Variance blocks. (c) An example of prediction errors and prediction updates across trials for one participant.

### 3. Results

#### 3.1 Prediction error weighting

In each trial, the *prediction error* is the spatial discrepancy between the participant's prediction and the target marker that appears after they make their prediction (ranging from -100% to 100% of the horizontal line). The *prediction update* is the spatial discrepancy between the prediction in a given trial and the prediction in the following trial. An example of prediction errors and prediction updates across the task for one participant is shown in *Figure 1c*. The *learning rate* (or *weighting of the prediction error*) is the prediction update divided by the prediction error. The learning rate is a measure of the extent to which a prediction error informs the next prediction. For instance, if, in a given trial, the participant's prediction moves only a small fraction towards the most recently displayed target marker, this indicates a low weighting of prediction error.

The primary measure computed for each participant was the median learning rate across trials. The median was used rather than the mean due to the occurrence of extreme scores in the learning rate data. Occasional extreme scores are expected because the learning rate is a ratio between prediction updates and predictions errors; for instance, an extreme score may result if the prediction error in a trial is close to zero but the participant changes their prediction nonetheless. It is possible for the learning rate in a given trial to be negative (indicating that the prediction was shifted *away* from the most recent target marker) or above 1 (indicating that prediction updates overshot the most recent target marker). In the present data, however, the median learning rate across the task was positive for all participants and typically below 1 ( $M = 0.65$ ;  $SD = 0.26$ ; range = 0.11–1.16). This indicates that prediction updates tended to be *towards* the most recent prediction error,

suggesting that participants used the prediction error feedback in each trial to update their predictions, as expected.

The median learning rate was lower in the second half of the task ( $M = 0.58, SD = 0.29$ ) compared to the first half of the task ( $M = 0.74, SD = 0.25$ ),  $F(1, 38) = 19.48, p < .001, \eta^2_p = .34$ . This is broadly consistent with Bayesian updating: as more data are sampled, the precision of prior beliefs increases and as a result the learning rate is reduced. There was no difference in this effect between participant groups, as indicated by a non-significant interaction effect,  $F(1, 38) = 0.04, p = .85, \eta^2_p < .01$ .

### **3.2 Block and Group comparisons**

A  $2 \times 2$  mixed model ANOVA was performed for median learning rate with Block as the within-subjects factor (High Variance vs. Low Variance) and Group as the between-subjects factor (High AQ vs. Low AQ).

A significant main effect of Block indicated that the learning rate was higher when the underlying Gaussian distribution that target markers were sampled from was more precise ( $M = 0.74, SD = 0.27$ ) compared to when it was less precise ( $M = 0.58, SD = 0.27$ ),  $F(1, 38) = 28.04, p < .001, \eta^2_p = .43$ . In other words, participants were more responsive to prediction errors when the incoming data was more precise. This is consistent with the Bayesian principle that the extent to which new sensory data should drive beliefs depends on its precision.

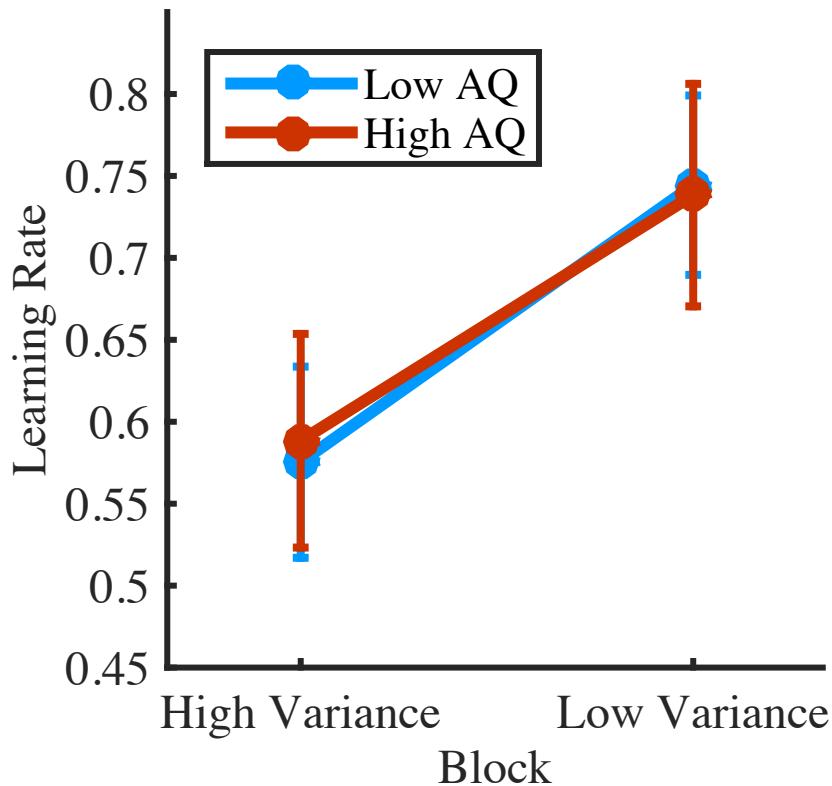
However, the main effect of Group was non-significant,  $F(1, 38) = 0.002, p = .97, \eta^2_p < .01$ . Similarly, the Block by Group interaction effect was non-significant,  $F(1, 38) = 0.10, p = .75, \eta^2_p < .01$ . Descriptive statistics are shown in *Figure 2*. These results indicate no support for a difference between groups in their weighting of prediction errors, conflicting with the hypothesis drawn from Pellicano and Burr (2012) and Van de Cruys et al. (2014) that autistic features are associated with a consistently higher weighting of prediction errors. Similarly, the lack of interaction effect indicates that learning rates were adjusted in response to the precision of target markers regardless of the individual's level of autistic traits, conflicting with

the notion that a chronically high learning rate linked to autistic features leads to an inflexibility in adjusting the weighting of prediction errors in line with the precision of sensory data.

A limitation of traditional significance testing is that a negative finding can indicate either that the null hypothesis is supported or that there is a lack of evidence to distinguish between the null and alternative hypotheses. A Bayesian analysis was therefore applied to establish whether the present data provide evidence specifically for the null hypothesis that High and Low AQ groups do not differ in their learning rate. Bayes factors reported here ( $\text{BF}_{01}$ ) quantify the evidence for the null hypothesis against the evidence for the alternative hypothesis. Values that approach 1 indicate that the data provide limited support for one hypothesis over the other. Similar to an effect size, the Bayes factor is a continuous measure that is typically interpreted with respect to rule of thumb labels (e.g.,  $\text{BF}_{01} > 3$  suggests moderate evidence for the null hypothesis,  $\text{BF}_{01} > 10$  suggests strong evidence, etc.; see Jeffreys, 1961; Lee & Wagenmakers, 2013). Bayesian analyses were performed with the open source statistical software JASP, Version 0.7.1 (Love et al., 2015).

A Bayesian independent sample *t*-test compared median learning rate across the full task between the High and Low AQ groups. Given the paucity of relevant prior information in the literature, a wide prior distribution was set (Cauchy prior width = 1; i.e., an uninformative prior) such that the test outcome was primarily informed by the present data. The Bayes factor indicated moderate support for the null hypothesis over the alternative hypothesis,  $\text{BF}_{01} = 4.24$ . These data therefore provide evidence that there is no difference in learning rate between High and Low AQ groups.

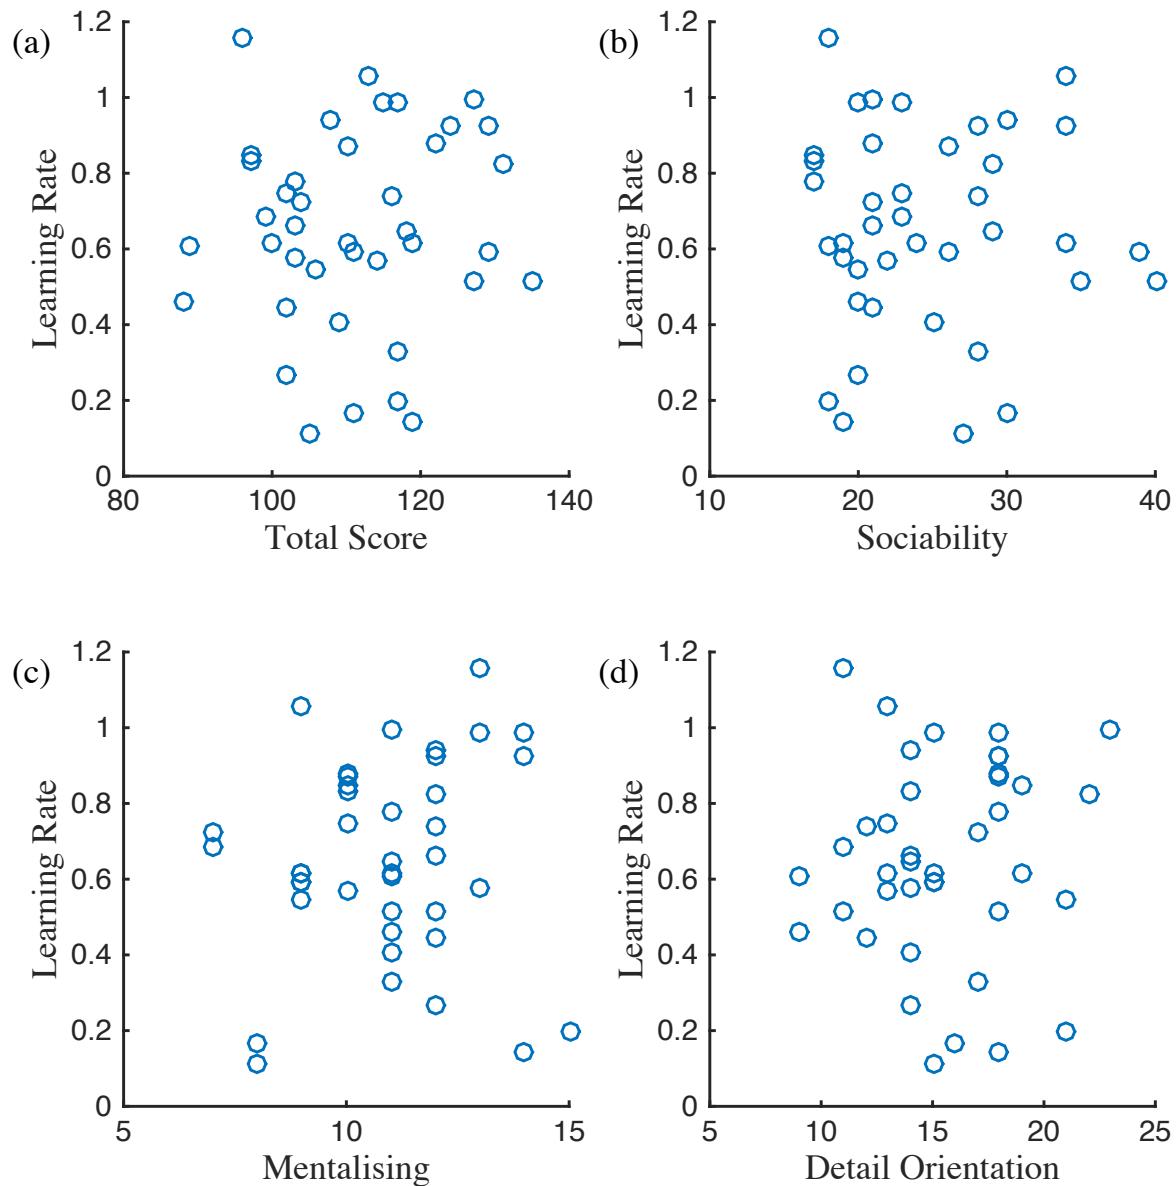
The same conclusions were supported when AQ score was included as a random factor in a Bayesian ANOVA with Block. This indicates that the lack of association between autistic traits and task performance reported here is not simply due to an effect of dichotomising the AQ variable via median split.



**Figure 2.** Mean learning rate across task blocks for groups defined by high and low autistic traits. Error bars indicate standard error.

### 3.3 Correlations

Pearson's correlations were performed to examine the extent to which median learning rate shared a linear relationship with total and subscale AQ scores across the full sample. AQ measures did not share any significant correlations with median learning rate ( $p > .05$ ), correlation coefficients were all very small ( $r = -.01-.09$ ) and Bayesian Pearson correlations (with Beta prior width = 1) indicated moderate support for the null hypothesis in each case ( $BF_{01} = 4.43-5.07$ ). These results are displayed in *Figure 3*, and indicate that there is no linear relationship between median learning rate and autistic traits.



**Figure 3.** There was evidence that median learning rate shared no relationship with autistic traits across the full sample. (a) The correlation between learning rate and total AQ score was non-significant,  $r = .07, p = .69$ , two-tailed,  $BF_{01} = 4.69$ . (b) The correlation between learning rate and Sociability was non-significant,  $r = -.01, p = .96$ , two-tailed,  $BF_{01} = 5.07$ . (c) The correlation between learning rate and Mentalising was non-significant,  $r = .09, p = .59$ , two-tailed,  $BF_{01} = 4.43$ . (d) The correlation between learning rate and Detail Orientation was non-significant,  $r = .06, p = .69$ , two-tailed,  $BF_{01} = 4.71$ .

## 4. Discussion

The present study examines inference in a statistical learning task and finds evidence *against* the hypothesis that autistic features in a non-clinical sample are associated with a persistently greater weighting of prediction errors. This conflicts with specific proposals for data-driven inference in ASD, as well as broader hypotheses regarding atypical precision weighting (e.g., Hohwy, 2013; Lawson et al., 2014; Pellicano & Burr, 2012; Van de Cruys et al., 2014). There are several implications of these negative findings that may be useful for developing Bayesian accounts of ASD.

First, this study tests the relatively simple hypothesis for a *chronic* difference in learning rate. To understand autistic cognition in Bayesian terms, we may need to appeal to more nuanced aspects of precision-weighted inference. One example is the role of *volatility processing*; Christoph Mathys and colleagues have developed a model of Bayesian inference in which the learning rate is flexibly modulated in part by inferences regarding the volatility of the environment – that is, the tendency for the hidden causes of sensory data to change over time (Mathys et al., 2011; Mathys et al., 2014). Thus, one possibility is that ASD is characterised by a deficit in volatility processing rather than a chronic difference in learning rate *per se*. This would be suggestive of a more context-dependent abnormality in learning rate in autistic samples rather than a persistently high learning rate (discussed in Palmer, Paton, Kirkovski, et al., 2015). Recent evidence in favour of this hypothesis comes from a study that examined learning of reward probabilities associated with choice options in a decision-making task (Robic et al., 2015); adults with ASD were similar to controls in learning appropriate choice behaviour when reward probabilities were stable over time, but performed significantly more poorly in a volatile context in which the underlying reward structure fluctuated.

Second, Bayesian inference in the brain is most commonly considered as a functional description of the *unconscious sensory processing* that underlies perception. In contrast, the task used in the present study entails *explicit* predictions and prediction errors, and thus can be considered more cognitive in

nature. The implication is that differences in learning rate that seem to capture important aspects of autistic perception may exist specifically for lower-level sensory processing rather than being relevant to more explicit forms of inference or statistical learning. There is considerable interest in the neuroscientific literature regarding more explicit forms of statistical learning processed in dorsomedial regions of frontal cortex (e.g., Behrens et al., 2007; McGuire et al., 2014); thus, the results of the present study help to address the natural question of whether differences in sensory processing hypothesized for ASD are characteristic of more cognitive as well as perceptual functions. This type of consideration is important partly because very similar differences in precision-weighting have been proposed to account for conditions other than ASD, such as schizophrenia (Adams et al., 2012; Adams, Stephan, et al., 2013; Edwards et al., 2012; Fletcher & Frith, 2009; Friston et al., 2014); one way of developing more specific predictive processing accounts of different psychiatric and neurological conditions may be in terms of *where* in the brain (or in what functions) aberrant precision weighting is most pronounced.

A limitation of the present study is that statistical learning is examined with respect to non-clinical autistic traits rather than in a clinical sample. However, a wide range of studies have found that psychophysical, cognitive and neurophysiological characteristics observed in clinical individuals also vary across the general population with respect to autistic traits. This includes, for example, aspects of weak central coherence (Grinter, Maybery, et al., 2009; Grinter, Van Beek, et al., 2009), susceptibility to sensory illusions (Palmer, Paton, Kirkovski, et al., 2015; Walter et al., 2009), sensory discrimination (Stewart, Griffiths, & Grube, 2015) and sensory hyper- and hypo-sensitivities (Robertson & Simmons, 2013). Bayesian theories account for each of these characteristics in the clinical condition with appeal to differences in the learning rate; for this reason, we should expect differences in Bayesian mechanisms that underlie autistic symptoms to similarly underlie individual variation in autistic traits in the general population. The sample size and range of AQ scores in the present study are comparable to previous studies that have reported correlations and between-group differences in perceptual and behavioural measures related to autistic traits (e.g., Palmer, Paton,

Enticott, et al., 2015; Palmer, Paton, Kirkovski, et al., 2015; Skewes et al., 2014; van Boxtel & Lu, 2013a).

In summary, these data provide evidence against the hypothesis that autistic features are associated with a persistently higher weighting of prediction errors, operationalized in the context of statistical learning. This may be important for the theoretical development of Bayesian accounts of ASD, indicating in particular that a *persistent* difference in learning rate may not be characteristic of the autistic phenotype (suggesting, perhaps, a more context-dependent atypicality in learning rate) or that differences in sensory processing that underlie perceptual features of ASD do not extend to cognitive processes underlying more explicit inference.

## Acknowledgements

The authors wish to thank Dror Cohen and Jochem Van Kempen for their assistance in data collection, and Tony Kemp for programming the task used in this experiment. This work was funded by an Australian Research Council Discovery grant (DP1311336). The authors have no conflict of interest to declare.

## Linking text between chapters 4 and 5

The empirical work reported in *Chapters 2, 3 and 4* examined aspects of perception, action and statistical learning partly with respect to variation in *autism-like traits* in the nonclinical population. This approach is motivated by evidence that sensory, cognitive and behavioural aspects of ASD represent an extreme end of a *spectrum* of characteristics that encompasses the population at large (reviewed in *Chapter 5*). An implication of this evidence is that the cognitive and neural mechanisms that underlie autistic symptoms may also vary in their function to a significant extent across the general population, accounting for individual differences more broadly. Examining perceptual and cognitive differences along the broader autism spectrum is thus an interesting testing ground for understanding how our experience of the world is produced by brain mechanisms. This in turn can help us to understand the nature of clinical conditions, such as ASD, in which sensory and cognitive differences impact strikingly upon the individual's quality of life. For instance, by examining the clinical symptoms of ASD together with non-clinical variation in autistic traits (as reported for body perception in *Chapter 3*), we can gain a broader understanding of how clinical individuals both differ and share similarities with different sections of the non-clinical population.

As noted in the *Preface*, the diagnosis of ASD is made on the basis of a diverse set of behaviours, such as reduced social responsiveness, strong preferences for routine, and sensory hypersensitivity. Individuals who receive this diagnosis show substantial *heterogeneity* in the way these characteristic symptoms present, as well as differing in whether *associated features* (such as intellectual disability) are present and in the underlying *etiological pathways* (e.g., multiple genetic causes; Geschwind & State, 2015; Newschaffer et al., 2002; Silberman, 2015). Correspondingly, there is ongoing debate regarding the extent to which ASD is a *unitary* condition, a set of *subgroups*, or comprised of *individual symptom domains* that are most appropriate to investigate independently from one another (e.g., Anderson, Montazeri, & de Bildt, 2015; Happé et al., 2006; Mandy & Skuse, 2008).

These general questions regarding how the ASD diagnosis is best conceptualised have implications for how we examine variation in autistic traits in the general population. The majority of research to date that investigates cognitive and neurophysiological mechanisms associated with autistic traits in adults treat autistic traits as a singular spectrum along which individuals differ primarily in severity alone: each individual is rated as being more or less autistic in general, and this total measure of autistic tendency is examined with respect to cognitive or neurophysiological measures of interest. An alternative approach, however, is to consider individual symptom domains independently: thus individuals may differ more meaningfully in their *pattern* or *profile* of traits across the different symptom domains of ASD.

To further illuminate the neurocognitive basis of the autism spectrum, an improved understanding of how autism-like traits manifest in adults will be valuable. This is in part a methodological point concerning whether empirical research would better proceed in examining autistic traits in adults as a singular dimension or as a set of independent trait dimensions. It is also relevant to the development of neurocognitive theories, however; for instance, there is a question of whether a core difference in information processing mechanisms in the brain is likely to underlie the myriad symptoms of ASD, or whether certain symptom domains are more likely to have distinct etiologies. In *Chapter 5*, we contribute to research in this area by analysing how autistic traits present in a large sample of adults recruited from the general community.

## Declaration for Thesis Chapter 5

### Declaration by candidate

In the case of *Chapter 5*, the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
<ul style="list-style-type: none"> <li>• Contributed to the experiment design.</li> <li>• Performed data analysis.</li> <li>• Contributed to the interpretation of results.</li> <li>• Wrote the paper.</li> </ul>	70%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Bryan Paton	Contributed to the experiment design, data analysis, interpretation of results, and writing.	15%
Peter Enticott	Contributed to the experiment design, interpretation of results, and writing.	N/A
Jakob Hohwy	Contributed to the experiment design, interpretation of results, and writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date 16/02/2016
--------------------------	---	--------------------

Main Supervisor's Signature		Date 16/02/2016
-----------------------------------	---	--------------------

# Chapter 5

## ‘Subtypes’ in the presentation of autistic traits in the general adult population

Colin J. Palmer<sup>1</sup>, Bryan Paton<sup>1, 2, 3</sup>, Peter G. Enticott<sup>4</sup>, Jakob Hohwy<sup>1</sup>

<sup>1</sup>*Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup>*Monash Biomedical Imaging, Monash University, Melbourne, Australia*

<sup>3</sup>*School of Psychological Sciences, Monash University, Melbourne, Australia*

<sup>4</sup>*Cognitive Neuroscience Unit, School of Psychology, Deakin University, Melbourne, Australia*

Published in *Journal of Autism and Developmental Disorders*, 2015

The final publication is available at Springer via <http://dx.doi.org/>

10.1007/s10803-014-2289-1

## Abstract

The present study examined the presentation of autistic traits in a large adult population sample ( $n = 2343$ ). Cluster analysis indicated two subgroups with clearly distinguishable trait profiles. One group ( $n = 1059$ ) reported greater social difficulties and lower detail orientation, while the second group ( $n = 1284$ ) reported lesser social difficulties and greater detail orientation. We also report a three-factor solution for the Autism-Spectrum Quotient, with two, related, social-themed factors (Sociability and Mentalising) and a third non-social factor that varied independently (Detail Orientation). These results indicate that different profiles of autistic characteristics tend to occur in the adult nonclinical population. Research into nonclinical variance in autistic features may benefit by considering social- and detail-related trait domains independently.

## 1. Introduction

A central feature of autism spectrum disorder (ASD) is the range in symptom severity: both within the *clinical* population, contributing to heterogeneity amongst individuals with ASD, and, of particular interest here, encompassing the *general* population (American Psychiatric Association, 2013; Lai, Lombardo, Chakrabarti, & Baron-Cohen, 2013). This latter impression comes primarily from research that has recorded the prevalence of autistic features across large community samples. Significantly, the extent to which the relevant traits are reported in the community appears to vary continuously and with normal or skewed-unimodal distribution (e.g., Baron-Cohen et al., 2001; Constantino & Todd, 2003; Hurst et al., 2007; Posserud et al., 2006). This contrasts, for example, with discontinuity or bimodality in trait distribution, either of which could set the clustering of symptoms in ASD apart as a matter of type rather than degree. Moreover, there is evidence for overlap in the genetic bases of autistic traits in the nonclinical population and ASD symptoms in the clinical population (Bolton et al., 1994; Chakrabarti et al., 2009; Piven, Palmer, Jacobi, Childress, & Arndt, 1997; Robinson et al., 2011).

A recent surge in cognitive and neuroscientific research builds on these findings by comparing nonclinical adults in their degree of autistic traits as an approach to studying the mechanisms underlying this condition. This has included studies examining visual search (e.g., Grinter, Maybery, et al., 2009; Grinter, Van Beek, et al., 2009), visual illusion susceptibility (Walter et al., 2009), biological motion processing (van Boxtel & Lu, 2013a), multisensory integration (Donohue et al., 2012; Palmer et al., 2013), social attention (Freeth, Foulsham, & Kingstone, 2013; Nummenmaa et al., 2012), gaze responses (Bayliss & Tipper, 2005; Chen & Yoon, 2011) and neurophysiological characteristics (Jakab et al., 2013; Nummenmaa et al., 2012; Puzzo et al., 2010; Sutherland & Crewther, 2010). These studies most commonly use the Autism-Spectrum Quotient (AQ), an inventory measure designed for use with adult community samples as either a screening or research tool (Baron-Cohen et al., 2001). The AQ necessarily assesses a range of social (e.g., sociability, mentalising) and non-social (e.g., restricted interests, detail-focus)

attributes, and researchers have typically compared individuals on their total AQ score or subscale scores, each produced by summing across items that assess more specific behaviours or characteristics.

It is an open question, however, regarding the extent to which ASD can be conceived of as a *singular* construct that varies across the general population. A ‘unitary spectrum’ model such as this conflicts with a suggestion by Happé, Ronald and colleagues that distinct genetic, neurological and cognitive causes may underlie the symptoms of ASD (Happé & Ronald, 2008; Happé et al., 2006; see also Mandy & Skuse, 2008). This alternate approach has been labelled the ‘fractionable autism triad’ hypothesis – in reference to the three symptom domains of impaired social interaction, impaired communication and restricted behaviours/interests specified for Autistic disorder in the DSM-IV (American Psychiatric Association, 2000) and ICD-10 (World Health Organization, 2004). (This model is still relevant to, and in some respects better represented by, the diagnostic criteria set out for ASD in the current 5th edition of the DSM, however; American Psychiatric Association, 2013)

The evidence that separate etiologies underlie the three DSM-IV symptom domains comes most notably from research into the relationship between autistic characteristics in the community, employing large-scale childhood twin samples (Robinson et al., 2012; Ronald, Happé, Bolton, et al., 2006; Ronald, Happé, & Plomin, 2005, 2006; Ronald, Larsson, Anckarsater, & Lichtenstein, 2011). While symptoms in the three domains co-occur at a level above chance (Happé & Ronald, 2008; see Mandy & Skuse, 2008 for review), low to moderate (rather than strong) correlations are observed between these characteristics. In addition, similarly modest overlap is estimated between the heritability of each domain, and, finally, individual symptoms can sometimes occur strongly in the absence of other cardinal features of ASD. A recent review of factor analytic studies involving clinical ASD participants concluded support for two distinct social and non-social dimensions underlying inventory measures of ASD symptoms – again suggesting that a conception of ASD as a unitary construct is insufficient (Shuster, Perry,

Bebko, & Toplak, 2013). Multi-factor solutions are similarly reported for nonclinical samples (e.g., Austin, 2005).

The idea here is that ASD trait domains vary independently to a marked degree across the population – with a diagnosis of ASD applied in cases where these characteristics happen to co-occur strongly. In principle, two individuals with the same AQ score, for instance, may still have a different *profile* of autistic traits, with corresponding cognitive and neurological dissimilarities. The present study aimed to build on the discussed literature by examining how profiles of autistic characteristics, as assessed by the AQ, tend to manifest in the general adult population. We used cluster analysis to examine the presentation of AQ item scores in a large sample recruited online. This statistical method finds the optimal manner of grouping cases (here, individuals) based on their similarity across a range of variables (here, the items of the AQ); in other words, this technique allowed us to discover patterns of trait presentation that tend to occur in the population. We also performed a factor analysis (as has been done previously for the AQ; e.g., Austin, 2005) to provide a more complete picture of the relationship between constructs underlying the AQ and how these constructs differ between individuals.

To our knowledge, this is the first study to look specifically at nonclinical variation in autistic traits in adults with cluster analysis, and one of the first to apply this technique to the AQ. The presentation of autistic traits in the adult population is important to clarify, due in part to the developmental nature of ASD and the present trend in research that focuses on nonclinical adult samples. Cluster analysis has been performed previously using the AQ with a high-functioning adult clinical group: Ring, Woodbury-Smith, Watson, Wheelwright, and Baron-Cohen (2008) report that their sample comprised between two and four clusters of individuals that varied in the severity of symptoms but not in their symptom profile. These authors note that this clustering supports a unitary spectrum (or ‘severity gradient’) characterisation of the heterogeneity in this clinical group, such that individuals are best grouped according to their overall magnitude of traits rather than in their profile of characteristics. Similarly, we can hypothesise that, if

a unitary spectrum account best describes the presentation of autistic characteristics in the general adult population, cluster analysis will reveal either a single cluster (i.e., no subgroups) or multiple clusters that differ in mean item or domain scores, but not in their profile of responses across the different traits that the AQ assesses. In contrast, the fractionable triad hypothesis allows for the possibility that the presentation of autistic traits in the general population is best characterised in terms of subgroups with different trait profiles.

## 2. Methods

### 2.1 Participants

The sample comprised 2343 adults who completed the AQ via an online survey platform ( $M = 31.36$ ,  $SD = 11.17$  years; 47.8% female; 86.6% right handers, 10.7% left handers; 88.0% indicating some degree of tertiary education). A link to the study was displayed in the Amazon Mechanical Turk system, and participation was restricted to individuals who identified as over 18 years in age and with current residence in the United States. Forty individuals who did not complete the survey and 63 individuals who shared an IP address with a previous participant were excluded prior to data analysis. Each participant received 1.50 USD upon completion of the study. Ethical approval for the study was obtained from the Monash University Human Research Ethics Committee. All participants gave informed consent.

### 2.2 AQ

The AQ is a 50-item self-administered questionnaire designed to assess both cardinal and associated traits of ASD in either clinical or nonclinical populations (Baron-Cohen et al., 2001). Each item consists of a statement (e.g., "When I'm reading a story, I find it difficult to work out the characters' intentions") that respondents rate in terms of personal applicability on a 4-point Likert scale ("definitely agree", "slightly agree", "slightly disagree", "definitely disagree"). In the

present study, we used Likert scoring (4-3-2-1) rather than the more common binary scoring method (1-1-0-0), as the former may better capture individual variability in responses and is more suitable for the analysis techniques employed here. Scores using this method can range 50–200 overall, or 10–40 for each of five subscales. Subscales are theoretically defined, consist of 10 items each, and are labelled Social Skills, Imagination, Communication, Attention to Detail and Attention Switching. Twenty-six items in the AQ are reverse scored, such that higher scores indicate greater resemblance to the symptoms of ASD for all items. Individuals with a diagnosis of ASD score significantly higher than nonclinical samples on both total AQ and individual subscales (Baron-Cohen et al., 2001; Broadbent, Galic, & Stokes, 2013; Hoekstra et al., 2008; Lau, Kelly, & Peterson, 2013). The test-retest reliability and internal consistency of this scale have each been evaluated as within an acceptable range (Baron-Cohen et al., 2001; Broadbent et al., 2013; Hoekstra et al., 2008; Lau, Kelly, et al., 2013).

### **2.3 Factor analysis**

An exploratory factor analysis (EFA) was conducted across the 50 items of the AQ using the maximum likelihood method of factor extraction. EFA employs a common factor analysis model as distinct from Principle Components Analysis. Promax (oblique) rotation was applied, bearing in mind that the traits that this questionnaire is designed to assess are expected to covary to an extent (Austin, 2005). EFA was performed from 1 to 10 maximum factors, and the factor solution that minimised the Akaike information criterion (AIC) was selected for further examination. AIC is preferred to significance testing here because the former weights the log likelihood evidence for each factor solution by the number of factors (model complexity). The scree plot, Kaiser's criterion, and parallel analysis were also examined for agreement with the AIC method in the number of factors to extract. The scree plot involves plotting the eigenvalues for each factor and retaining factors to the left of the point of inflexion. Kaiser's criterion involves retaining factors with eigenvalues greater than 1 (these methods, including their potential deficiencies, are reviewed in Field, 2009; Ledesma & Valero-Mora, 2007). Parallel analysis compares eigenvalues of each factor against the 95<sup>th</sup> percentile of

eigenvalues obtained from a set of control samples. In the present application, we used 1000 control samples that were permutations of the original data. The permutation method and other details are described further in O'Connor (2000, 2014). Factor scores for each participant were derived using ridge regression. Confidence intervals (95%) for item factor loadings were obtained via exhaustive delete-one jackknife resampling. Items that loaded on a given factor with a lower confidence bound greater than 0.4 were examined for factor interpretation, as this indicates at least 16% shared variance between the item and factor (Stevens, 2009). Cronbach's unstandardized alpha was calculated to assess internal consistency of the entire questionnaire and item subsets identified for each factor.

## 2.4 Cluster analysis

Clustering of the sample was investigated with cases characterised by responses across the 50 items of the AQ. Proximity between cases was quantified in terms of Spearman's rank correlation coefficient, and clusters were formed via hierarchical agglomerative clustering with complete linkage (consistent with Ring et al., 2008). Martinez, Martinez, and Solka (2011) recommends examining data clustering with the use of several distance and clustering methods: in the current dataset, the number of clusters observed did not differ when instead employing the Euclidean distance dissimilarity measure or Ward's and weighted average linkage methods.

The gap statistic method was used to determine whether or not clustering was apparent in the data and to estimate the number of clusters present (Martinez et al., 2011; Tibshirani, Walther, & Hastie, 2001). Participant data was first clustered separately for 1–10 clusters such that the within-cluster dispersion for each of these clustering outcomes could be calculated. Reference distributions were then generated as uniform distributions across the range of participant data for each variable. For each number of clusters (1–10), 10 reference distributions were created and underwent clustering. The average within-cluster dispersion of the 10 reference distributions for each number of clusters was calculated to form the dispersion expected if no real clustering was present in the data. The gap statistic for each cluster number is the difference between the observed and expected

dispersions. The number of clusters in the participant data is taken as the smallest number of clusters that satisfies that criteria  $gap(k) \geq gap(k+1) - s_{k+1}$ , where  $gap$  is the gap statistic,  $k$  is the number of clusters, and  $s$  the (weighted) standard deviation of the within-cluster dispersion of the 10 reference distributions. This procedure followed Tibshirani et al. (2001) and Martinez et al. (2011), and was run using MATLAB code adapted from the latter.

Gaussian mixture modelling was also applied to examine whether the number of clusters observed using the above techniques could be reproduced using a different method. This technique employs a generative variational Bayesian modelling procedure to find a set of Gaussian densities that best predict the observed data. It minimises a free energy quantity to determine the optimal number of clusters. This involves maximizing the negative log evidence for the model given the data. This procedure was run using MATLAB code provided by Daunizeau, Adam, and Rigoux (2014).

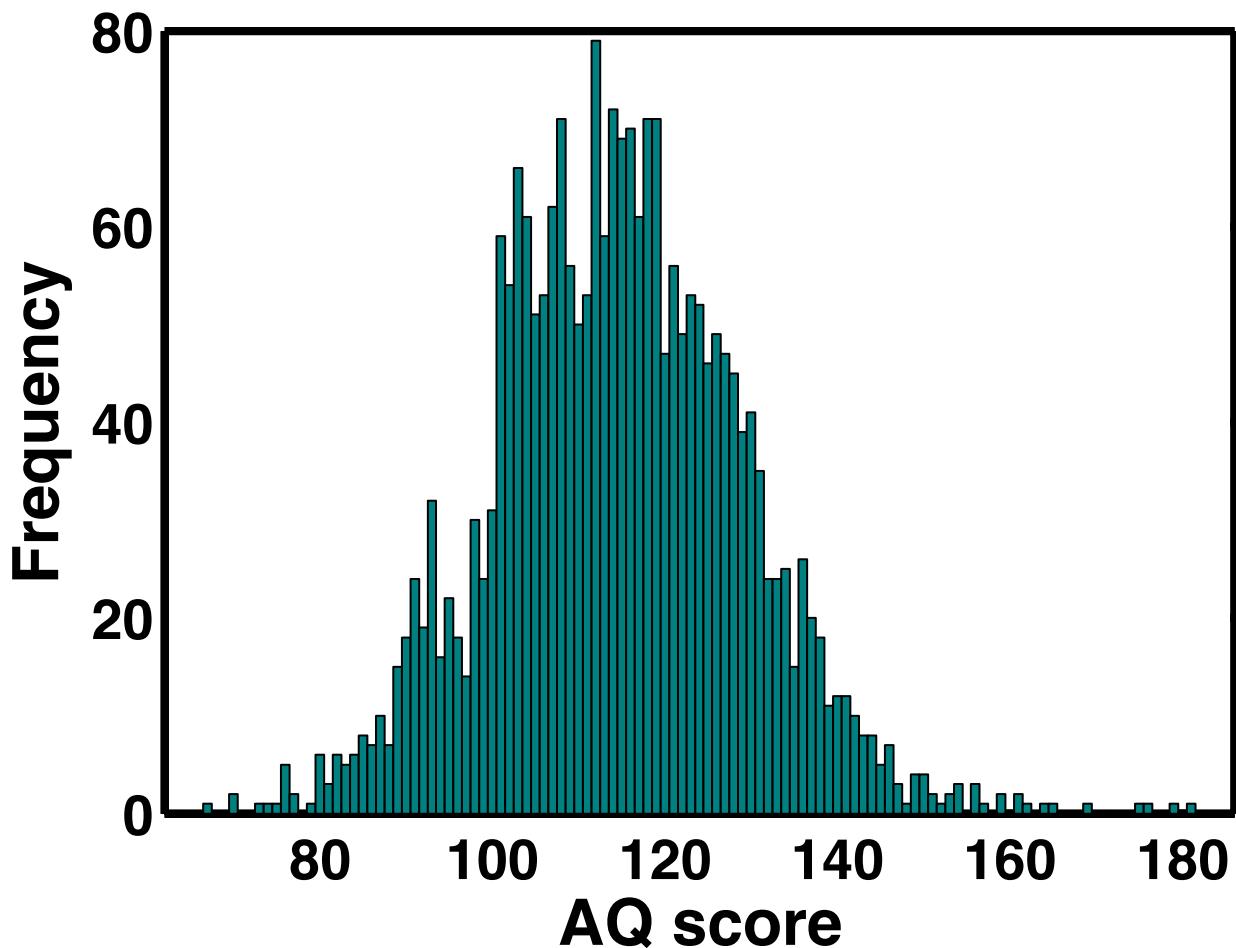
To characterise the observed clusters, mean scores for each item of the AQ were plotted for each cluster. Confidence intervals (95%) of item mean scores were computed using the jackknife resampling method. The proportion of items that matched a consistent severity gradient across clusters was examined for each subscale. Total AQ score, subscale scores, factor scores, and demographic characteristics (gender, age, handedness, and education level) were also compared between clusters using independent samples  $t$ -tests or Chi-square tests for independence (with Yates Continuity Correction for 2x2 contingency tables).

Data analysis was conducted in MATLAB, version R2012a (7.14.0.739).

### 3. Results

#### 3.1 AQ score distribution

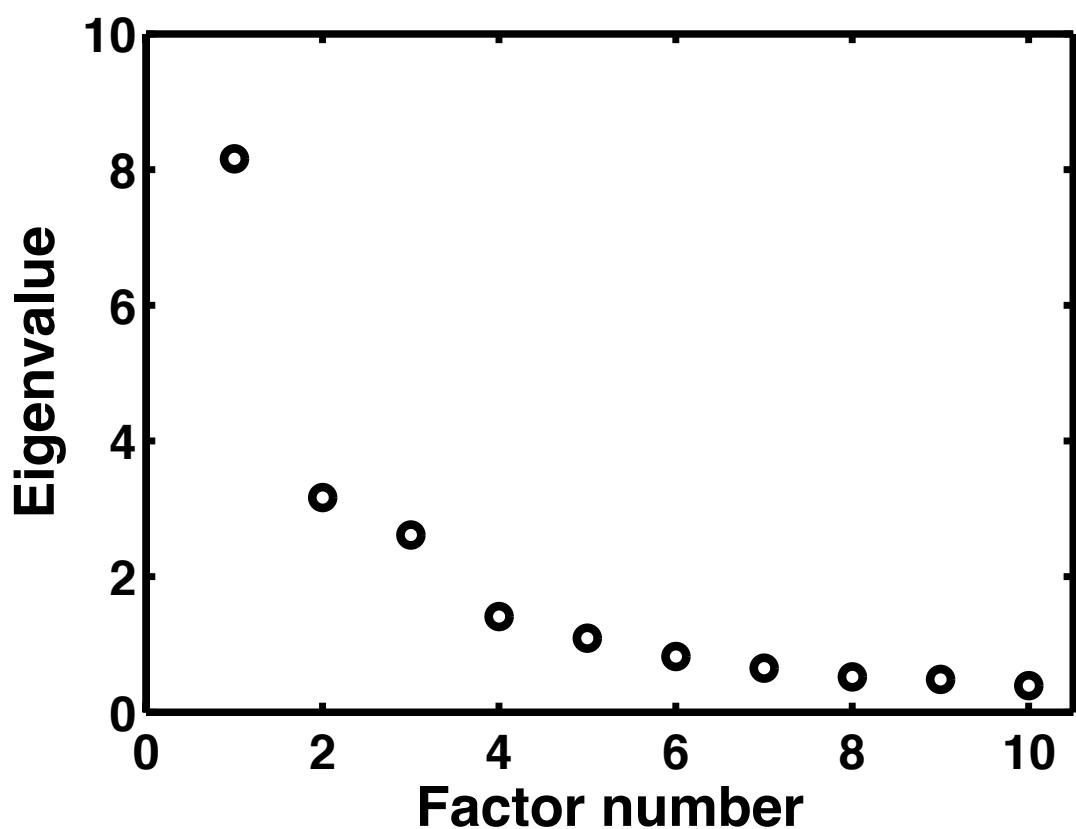
The mean total AQ score was 114 ( $SD = 14.5$ ; median = 114; median absolute deviation = 10), corresponding to a mean binary total score of 20 ( $SD = 6.9$ ; median = 20; median absolute deviation = 5). The distribution of total AQ scores is shown in *Figure 1*.



**Figure 1.** Distribution of total AQ scores across the full sample ( $n = 2343$ ).

### 3.2 Factor analysis

AIC values and variance accounted for by each factor solution are shown in *Table 1*. The factor solution that minimised the AIC comprised 3 factors and accounted for 27% of the variance. The scree plot, shown in *Figure 2*, also supported extraction of 3 factors. In contrast, Kaiser's criterion indicated 5 factors, while parallel analysis indicated 16 factors. Previous research using similar datasets has also noted over-extraction of factors when using parallel analysis (e.g., Stewart & Austin, 2009). The AIC method is preferable because it takes into account model evidence and model complexity. Thus, together there was strongest support for the three-factor solution. Items that loaded above threshold on each factor are shown in *Table 2*. A very clear pattern in item content can be seen such that the three factors correspond, respectively, to reduced sociability (reduced social skills or social discomfort), reduced mentalising (difficulty in mental-state attribution) and detail orientation.



**Figure 2.** Scree plot for exploratory factor analysis with common factor analysis model.

**Table 5.** AIC values and variance explained for factor solutions 1–10

Number of factors	AIC value	Variance accounted for
1	18.00	15.70
2	15.48	21.43
3	13.77	26.99
4	14.20	29.00
5	15.00	31.50
6	16.15	33.39
7	17.39	35.27
8	18.92	36.49
9	20.48	38.29
10	22.10	40.21

**Table 2.** Item loadings for three-factor solution.

Item content	Loading
<i>Factor 1: Sociability</i>	
I am good at social chit-chat. †	0.83
I find social situations easy. †	0.81
I enjoy social chit-chat. †	0.80
I enjoy social occasions. †	0.80
I enjoy meeting new people. †	0.79
I find it hard to make new friends.	0.63
I find myself drawn more strongly to people than to things. †	0.59
New situations make me anxious.	0.56
I would rather go to a library than a party.	0.55
I frequently find that I don't know how to keep a conversation going.	0.55
I prefer to do things with others rather than on my own. †	0.48
<i>Factor 2: Mentalising</i>	
I find it difficult to work out people's intentions.	0.63
When I'm reading a story, I find it difficult to work out the characters' intentions.	0.63
I find it easy to work out what someone is thinking or feeling just by looking at their face. †	0.53
I find it easy to "read between the lines" when someone is talking to me. †	0.51
When I'm reading a story, I can easily imagine what the characters might look like. †	0.50
I know how to tell if someone listening to me is getting bored. †	0.47
<i>Factor 3: Detail Orientation</i>	
I notice patterns in things all the time.	0.63
I usually notice car number plates or similar strings of information.	0.59
I am fascinated by numbers.	0.57
I am fascinated by dates.	0.55
I tend to notice details that others do not.	0.53
I like to collect information about categories of things (e.g. types of car, types of bird, types of train, types of plant, etc.).	0.51

Items with lower CI exceeding 0.4 loading threshold are shown. †-symbol indicates reverse-scored items.

The two social-themed factors shared a moderate positive association, but correlated very weakly in the negative direction with the Detail Orientation factor (*Table 3*). The associations between factor scores and the theoretically-defined subscales of the AQ (Baron-Cohen et al., 2001) demonstrated a similar pattern (*Table 4*). Specifically, the Sociability and Mentalising factors shared mostly moderate and strong positive correlations with the three social-themed subscales (Social Skills, Imagination, Communication) and the Attention-Switching subscale, but correlated negatively with the Attention to Detail subscale. In contrast, the Detail Orientation factor correlated strongly in the positive direction with the Attention to Detail subscale, but showed very weak correlations with the other four subscales. Similarly, the social subscales of the AQ tended to vary together moderately or strongly, but shared weak negative relationships with the Attention to Detail subscale (*Table 5*). The association between the Attention Switching and social domains of the AQ has been reported previously in the literature (Lau, Gau, et al., 2013; Murray, Booth, McKenzie, Kuenssberg, & O'Donnell, 2013; Stewart & Austin, 2009). This (perhaps counterintuitive) finding is discussed further in the Discussion.

**Table 3.** Factor intercorrelations

	Mentalising	Detail Orientation
Sociability	.48*	-.09*
Mentalising		-.14*

Pearson's linear correlation coefficients are shown. Bonferroni correction was applied to the alpha level.

\*  $p < .0001$  (two-tailed).

**Table 4.** Correlations between factor scores and AQ subscales

	Total Score	Social Skill	Imagination	Communication	Attention to Detail	Attention Switching
Sociability	0.79* [0.78, 0.81]	0.93* [0.93, 0.94]	0.27* [0.23, 0.30]	0.69* [0.67, 0.71]	-0.11* [-0.15, -0.07]	0.61* [0.58, 0.63]
Mentalising	0.72* [0.70, 0.74]	0.59* [0.56, 0.61]	0.68* [0.66, 0.70]	0.78* [0.76, 0.80]	-0.28* [-0.32, -0.24]	0.52* [0.49, 0.55]
Detail	0.28* [0.25, 0.32]	-0.08* [-0.12, -0.04]	0.00 [-0.04, 0.04]	0.11* [0.07, 0.15]	0.84* [0.83, 0.85]	0.04 [0.00, 0.08]
Orientation						

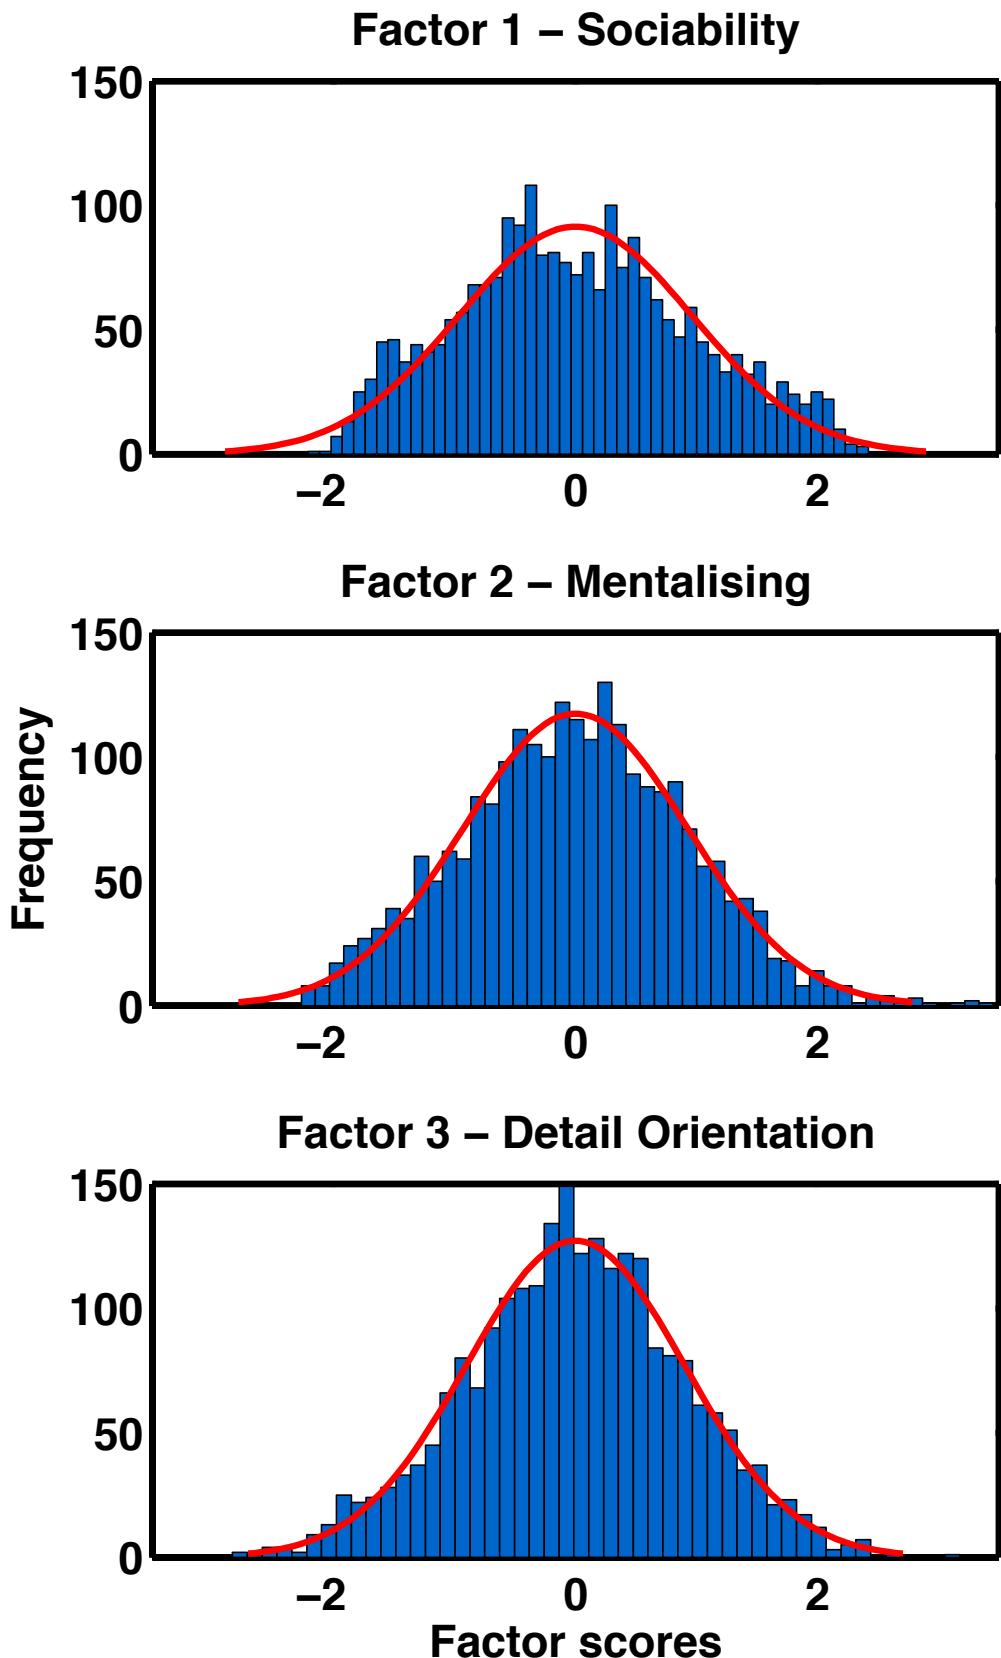
Pearson's linear correlation coefficients are shown with lower and upper bounds of the 95% confidence interval in square brackets. Bonferroni correction was applied to the alpha level. \*  $p < .0001$  (two-tailed).

**Table 5.** Intercorrelations between AQ subscales

	Social Skill	Imagination	Communication n	Attention to Detail	Attention Switching
Total Score	0.81* [0.80, 0.82]	0.58* [0.56, 0.61]	0.81* [0.80, 0.83]	0.20* [0.16, 0.24]	0.70* [0.68, 0.72]
Social Skill		0.32* [0.28, 0.35]	0.69* [0.67, 0.71]	-0.13* [-0.17, -0.09]	0.55* [0.52, 0.58]
Imagination			0.41* [0.37, 0.44]	-0.05 [-0.09, 0.00]	0.23* [0.19, 0.27]
Communication				-0.09* [-0.13, -0.05]	0.50* [0.47, 0.53]
Attention to Detail					-0.09* [-0.13, -0.05]

Pearson's linear correlation coefficients are shown with lower and upper bounds of the 95% confidence interval in square brackets. Bonferroni correction was applied to the alpha level. \*  $p < .0001$  (two-tailed).

The three factors identified each had high internal consistency ( $\alpha_{\text{Sociability}} = .90$ ,  $\alpha_{\text{mentalising}} = .74$ ,  $\alpha_{\text{Detail orientation}} = .75$ ), as did the full questionnaire ( $\alpha_{\text{AQ}} = .84$ ). The distribution of factor scores was approximately normal for each factor (*Figure 3*).

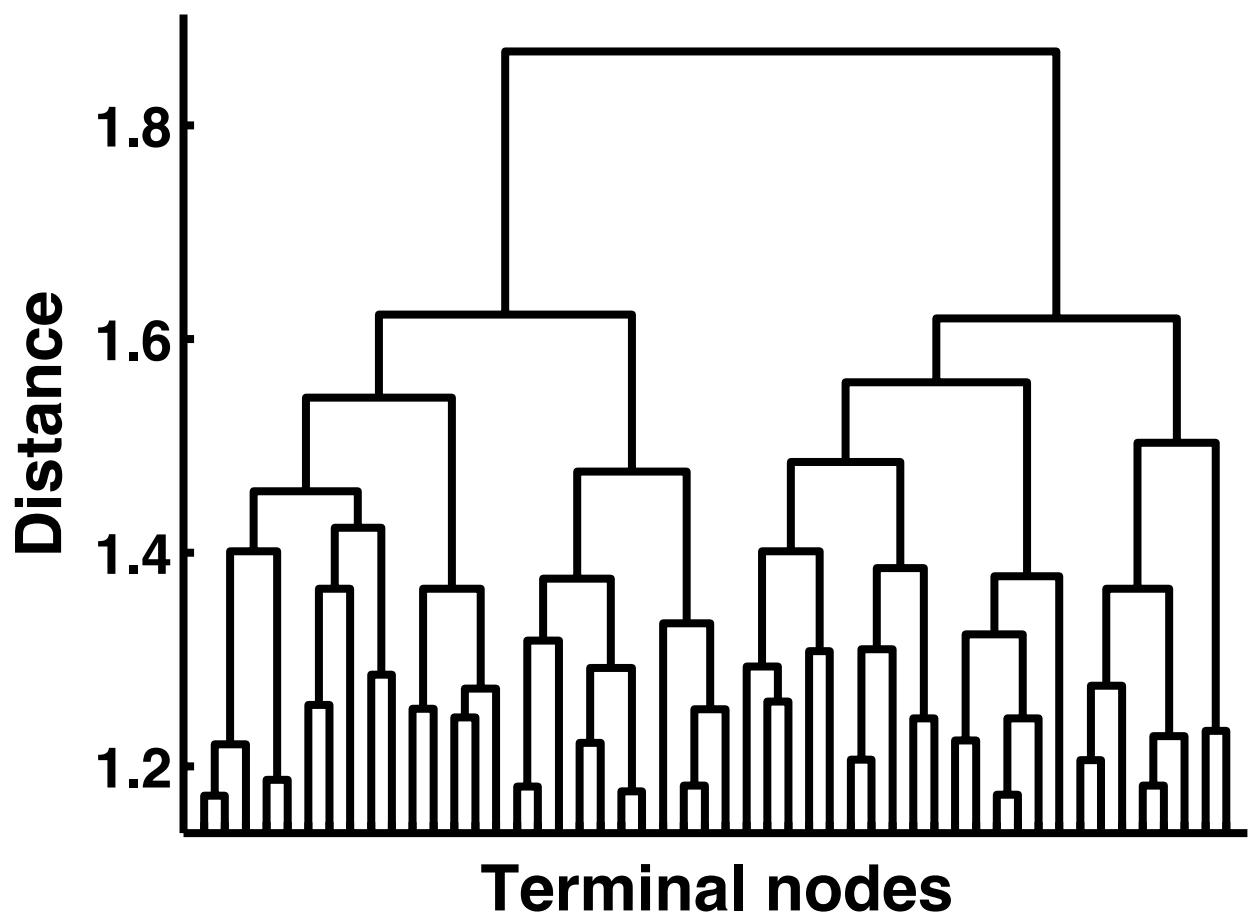


**Figure 3.** Distribution of factor scores across the full sample ( $n = 2343$ ). Gaussian distribution fits are shown.

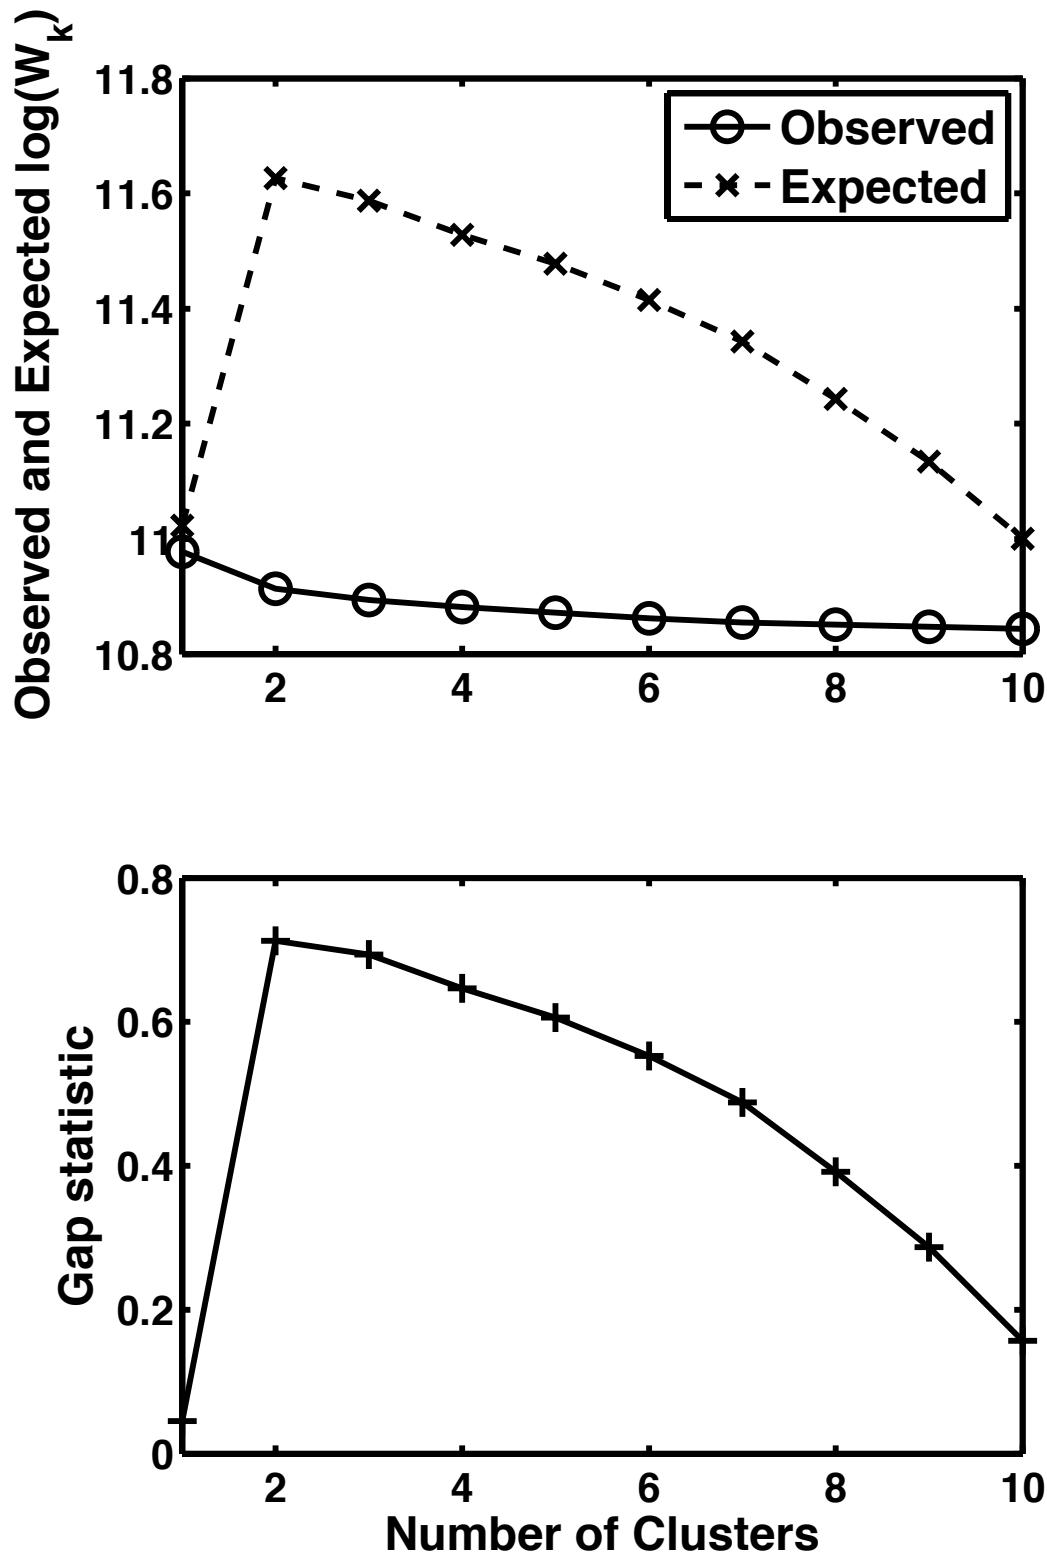
Females scored slightly higher on (reduced) Sociability than males, indicating that females were *less* sociable on average ( $M_{FEMALES} = 0.04$ ,  $SD_{FEMALES} = 1.02$ ;  $M_{MALES} = -0.04$ ,  $SD_{MALES} = 0.92$ ),  $t(2260) = -2.05$ ,  $p < .05$ ,  $r = .04$ . In contrast, males scored higher than females on (reduced) Mentalising ( $M_{MALES} = 0.12$ ,  $SD_{MALES} = 0.93$ ;  $M_{FEMALES} = -0.13$ ,  $SD_{FEMALES} = 0.92$ ),  $t(2341) = 6.45$ ,  $p < .001$ ,  $r = .13$ , and Detail Orientation ( $M_{MALES} = 0.09$ ,  $SD_{MALES} = 0.86$ ;  $M_{FEMALES} = -0.10$ ,  $SD_{FEMALES} = 0.94$ ),  $t(2275) = 4.93$ ,  $p < .001$ ,  $r = .10$ . Previous research has also tended to find that males score more highly in the AQ than females (e.g., Baron-Cohen et al., 2001); however, the effect sizes for the differences in factor scores between females and males in the present study are small.

### 3.3 Cluster analysis

The gap statistic procedure indicated two clusters of individuals within the sample (C1, C2; see *Figure 4* for dendrogram plot). *Figure 5* shows the observed dispersion, expected dispersion and gap statistic for 1–10 clusters. As can be seen, a single maximum occurs for the gap statistic at 2 clusters. The values of  $gap(k)$  -  $gap(k+1)$  -  $s_{k+1}$  for  $k = 1-9$  were: -0.63, 0.05, 0.09, 0.08, 0.09, 0.10, 0.13, 0.13, 0.13. The smallest number of clusters that satisfied the criteria  $gap(k) \geq gap(k+1) - s_{k+1}$  was two, indicating two clusters in the data. The two clusters were of roughly equal size ( $n_{C1} = 1059$ ;  $n_{C2} = 1284$ ). The Gaussian mixture modelling procedure indicated that a two-cluster model minimized free energy, providing supporting evidence that the correct number of clusters was obtained.



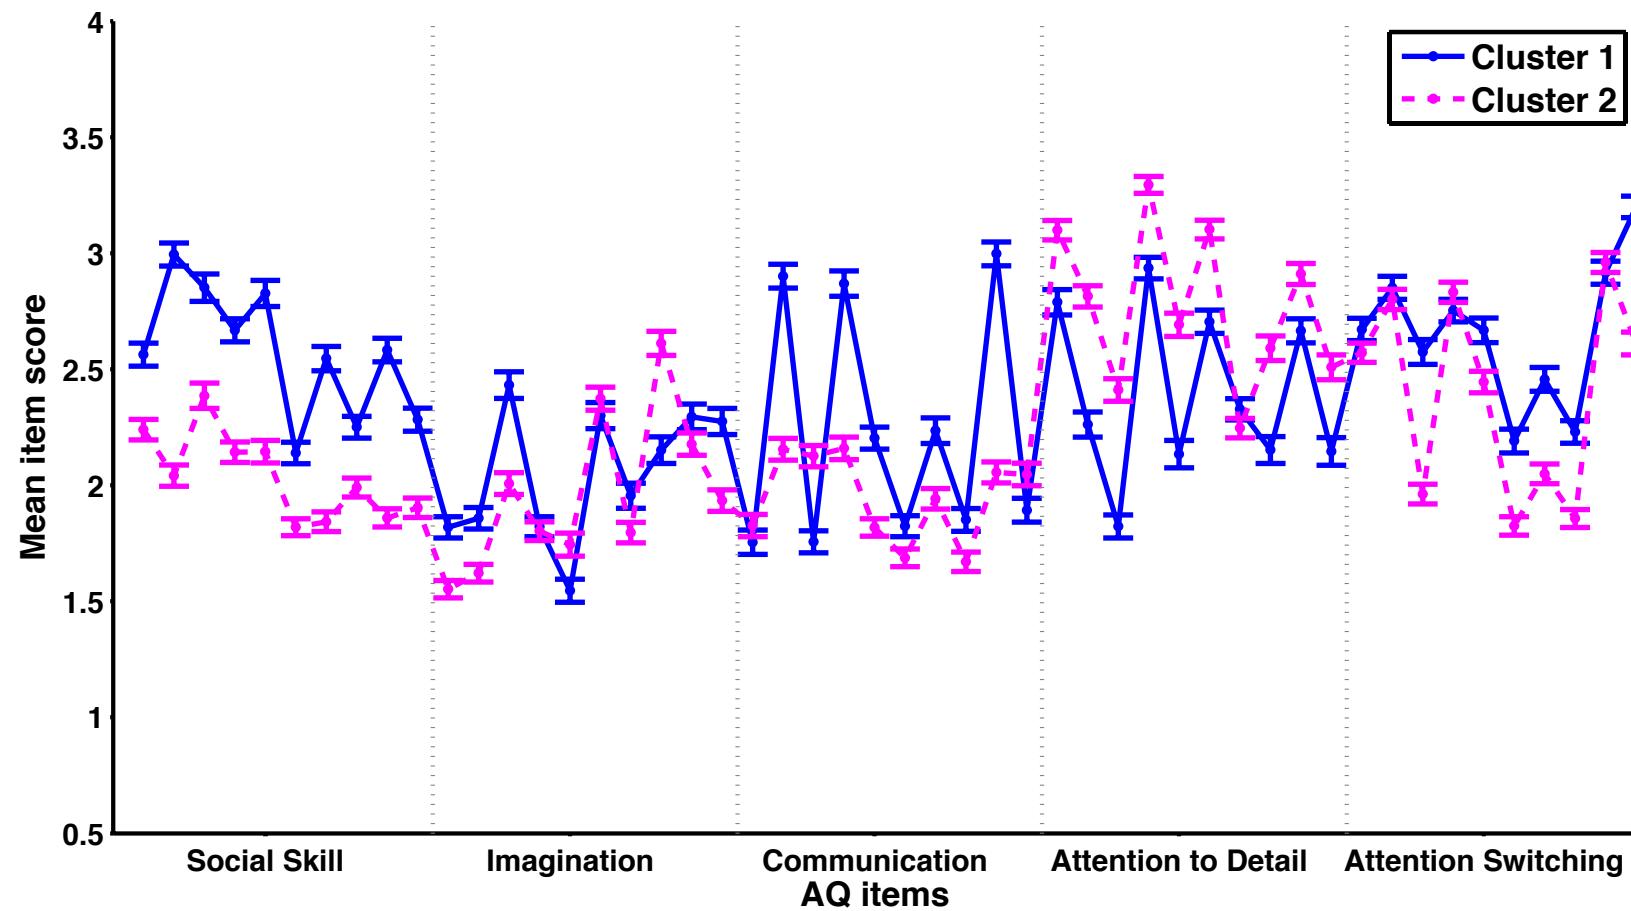
**Figure 4.** Dendrogram plot with 50 terminal nodes. Hierarchical agglomerative clustering with Spearman's proximity and complete linkage was applied to the sample with cases characterised by scores on the 50 items of the AQ.



**Figure 5.** The gap statistic procedure indicated two clusters within the sample. The upper plot shows observed and expected dispersion measures for clustering with 1–10 clusters. The lower plot shows the gap statistic for clustering with 1–10 clusters, with a maximum at two clusters.

There was a small, significant association between gender and cluster status,  $\chi^2 (1, n = 2343) = 21.10, p < .001$ , phi = -.10, with a greater proportion of females in C1 (53.1%) than C2 (43.5%). Individuals in C1 ( $M = 32.16, SD = 11.78$  years) were also marginally older on average than those in C2 ( $M = 30.69, SD = 10.60$  years),  $t(2152) = 3.15, p < .01, r = .07$ . There was a very weak association between education level and cluster status,  $\chi^2 (1, n = 2343) = 5.30, p < .05$ , phi = .05, with a slightly greater proportion of respondents indicating some degree of tertiary education in C2 (89.4%) than in C1 (86.2%). Finally, there was no significant association between handedness and cluster status,  $\chi^2 (2, n = 2343) = 2.04, p = .36$ , Cramer's V = .03 (C1: 86.9% right handers, 11.0% left handers; C2: 86.4% right handers, 10.5% left handers). In summary, the two clusters differed little in these four demographic characteristics, given the very weak effect sizes and the contribution of large sample sizes to the statistical significance of these comparisons. Thus, clusters were defined mainly in terms of autism-related features (as described below) rather than demographic characteristics.

The profile of mean scores across the 50 items of the AQ is shown for each cluster in *Figure 6*. As can be seen, C1 typically scored higher than C2 on items that comprise the three social subscales (C1 > C2 for all Social Skill items, 6 out of 10 Imagination items, and 7 out of 10 Communication items) and the Attention Switching subscale (C1 > C2 for 7 out of 10 items). In contrast, C2 scored higher than C1 for 9 out of 10 items that comprise the Attention to Detail subscale. The pattern of C1 scoring higher than C2 in social-themed characteristics (impairments) but lower in detail-themed characteristics is also apparent in the comparison of subscale scores and factor scores across clusters in *Table 6*. The effect sizes for comparisons of AQ subscale scores and factor scores between the two cluster groups were most commonly in the moderate range, as reported in *Table 6*. As can be seen, C1 also scored higher in total AQ than C2 (scoring higher on 60% items in the questionnaire).



**Figure 6.** Profile of mean AQ item scores for each cluster with 95% confidence intervals.

**Table 6.** Comparison of AQ subscale scores and factor scores across clusters

Variable	Cluster 1 <i>M (SD)</i>	Cluster 2 <i>M (SD)</i>	<i>t (df)</i>	<i>r</i>
<i>Cluster 1 mean &gt; Cluster 2 mean</i>				
Total AQ	118.88 (14.74)	111.05 (13.28)	13.39 (2151)	0.28*
Social Skill	25.70 (5.09)	20.37 (4.75)	26.02 (2192)	0.49*
Imagination	20.45 (4.25)	19.61 (3.83)	4.93 (2154)	0.11*
Communication	22.28 (4.40)	19.48 (4.35)	15.43 (2341)	0.30*
Attention	26.51 (4.23)	23.91 (3.84)	15.43 (2161)	0.32*
Switching				
Sociability	0.56 (0.85)	-0.46 (0.80)	29.99 (2199)	0.54*
Mentalising	0.29 (0.90)	-0.24 (0.88)	14.28 (2341)	0.28*
<i>Cluster 2 mean &gt; Cluster 1 mean</i>				
Attention to Detail	23.94 (4.45)	27.67 (4.05)	-21.02 (2162)	0.41*
Detail Orientation	-0.42 (0.86)	0.35 (0.78)	-22.35 (2163)	0.43*

Bonferroni correction was applied to the alpha level. \*  $p < .0001$  (two-tailed).

## 4. Discussion

The present study provides evidence for two patterns of subclinical autistic traits in the general adult population. In particular, cluster analysis indicated a group of respondents characterised by greater social difficulties and weaker detail-orientation, and a second group featuring better social abilities and stronger detail-orientation. This conflicts with a unitary spectrum model of autistic traits, which predicts either clustering based on symptom severity alone or an absence of subsets within the sample. (On the other hand, our results suggest that social characteristics – mentalising and sociability – *do* vary together to a moderate extent, and show a consistent severity gradient across the two subgroups identified). Clinical individuals score higher than controls on *all* subscales of the AQ (Baron-Cohen et al., 2001; Broadbent et al., 2013; Hoekstra et al., 2008; Lau, Kelly, et al., 2013), suggesting that the differences in profile between clusters reported in the present study cannot be explained as simply indicating that detail-orientation traits are extraneous to ASD. Rather, our findings accord with a view that social and detail-orientation characteristics related to ASD tend not to co-occur in the adult population – to the extent that nonclinical individuals are best grouped in terms of profile differences rather than in severity differences alone. Overall, this point sits well with the fractionable autism triad hypothesis, which suggests a degree of independent etiology for the three DSM-IV symptom domains for Autistic disorder (Happé & Ronald, 2008; Happé et al., 2006).

The results of the cluster analysis described here conflict with those of a comparable analysis employing the AQ in a high-functioning clinical sample ( $n = 333$ ; aged 16 years and above, including adults) (Ring et al., 2008). This prior study reports between two and four clusters distinguishable in terms of symptom severity but not in symptom profile. The discrepancy between studies could reflect a difference between the presentation of traits in clinical and nonclinical populations; however, it is important to appreciate that the former are *defined* by the co-occurrence of the diagnostic symptom domains, and thus Ring et al. (2008) address a question of how heterogeneity manifests when these symptoms co-occur, rather than how the relevant traits tend to present in general. Other

research that has examined the clustering of individuals within clinical ASD samples has typically focussed on children and employed measures other than the AQ. A recent latent class analysis of a large sample of nonclinical children, using questionnaire items based on DSM-IV diagnostic criteria, indicated four subgroups that differed not only in severity but also in symptom profile (Beuker et al., 2013). This latter finding is compatible with our results, although represents autistic characteristics at a significantly different stage of development.

The three-factor solution described in the present study indicates that the AQ is tapping into a 'sociability' construct, a 'mentalising' construct, and a 'detail orientation' construct. This finding closely matches those of two previous factor analyses of the AQ in smaller adult nonclinical samples (Austin, 2005; Hurst et al., 2007). Other researchers, however, have reported 2, 4, or 5 factors (Freeth, Sheppard, Ramachandran, & Milne, 2013; Hoekstra et al., 2008; Kloosterman, Keefer, Kelley, Summerfeldt, & Parker, 2011; Stewart & Austin, 2009). The three-factor solution provides empirical support for three item subsets in the AQ: one subset resembling the existing Attention to Detail subscale, and two other subsets characterised, respectively, by 'sociability' and 'mentalising' items drawn from the three traditional social subscales. These results do not attest to the five subscales typically used in research, however, which include separate Social Skills, Communication, Imagination and Attention-Switching subscales.

It is worth reiterating that the Attention Switching subscale correlated positively with the three social-themed subscales, as well as with the Sociability and Mentalising factors identified empirically in the present study. This subscale consists of items such as "I prefer to do things the same way over and over again", could be alternately labeled as 'behavioural inflexibility', and *a priori* would be considered a non-social feature of ASD. Similarly, there was no evidence from the factor analysis that the AQ items were tapping into a distinct 'attention switching' construct. Thus, findings based upon AQ scores (including the differences between clusters reported here) appear to be best explained in terms of social and detail-orientation characteristics. The positive association between Attention Switching items and social domains of the AQ may indicate that individual items in this

subscale tap into social characteristics and share this in common to a greater degree than their tendency to tap into a unitary construct of attention switching. Another possibility is that attention switching is a trait (or set of traits) that benefits social behaviour.

The pattern of moderate-to-strong positive correlation between the Sociability and Mentalising factors, while the third Detail Orientation factor varied more independently, coheres reasonably well with the DSM-5 criteria for ASD, which encompass a social-communicative domain and a second, non-social domain (American Psychiatric Association, 2013). A recent review of factor analytic studies employing clinical ASD samples similarly concludes that the most support exists for a two-factor solution consisting of distinct social and non-social domains (Shuster et al., 2013). The plurality of underlying constructs observed in the current and previous studies is consistent with the fractionable triad hypothesis, which suggests that the domains of autistic traits come apart to vary somewhat independently across the population. In the present study, a 'detail orientation' dimension stands out most distinctly from the rest of the AQ. Thus, the 'fractionation' in autistic traits in adults may occur most prominently between social (sociability and mentalising) and detail-orientation domains.

In terms of studying ASD via nonclinical samples, the present results indicate that adults tend to present with one of two differing combinations of traits. This implies that an individual lower on total AQ might actually be higher in detail orientation (and hence, more similar to a clinical individual on this particular symptom set) than an individual higher in total AQ. Thus, it is preferable to examine cognitive and neurological measures of interest with respect to subsets of the questionnaire rather than to total AQ score. Our results demonstrate two main domains that present independently, to an extent, in the general population – social characteristics and detail orientation – indicating that individuals are best characterised with regard to their autistic presentation with the use of these two dimensions. Further to this point, the majority of items in the AQ assess social features of ASD, which may be important to note for studies examining cognitive or neurological functions in relation to composite AQ scores in nonclinical samples.

(with detail-orientation arising here as the primary non-social feature that the AQ assesses). The use of more specific, and empirically-supported, item subsets may enhance the likelihood of relating nonclinical findings to the heterogeneous clinical population.

A conceivable limitation of the present study is the use of an online sample, recruited via the Amazon Mechanical Turk system. The demographics of this population are well-described in several prior investigations, however, and it has been argued that samples recruited in this fashion are more representative of the general (U.S.) population than tertiary student samples typically employed in psychological research (Berinsky, Huber, & Lenz, 2012; Buhrmester, Kwang, & Gosling, 2011; Paolacci, Chandler, & Ipeirotis, 2010). In the present study, the average score on the AQ ( $M = 20$ ,  $SD = 6.9$ ) was higher than that reported previously for student nonclinical samples (e.g.,  $M = 17.6$ ,  $SD = 6.4$ ) (Baron-Cohen et al., 2001), but not drastically so, and not approaching that reported previously for clinical samples (e.g.,  $M = 35.8$ ,  $SD = 6.5$ ) (Baron-Cohen et al., 2001). The Amazon Mechanical Turk and other online recruitment systems are increasingly commonplace in scientific research (e.g., Kidd & Castano, 2013; Quoidbach, Gilbert, & Wilson, 2013). The use of online recruitment in the present experiment allowed for the analysis of a notably larger sample than that employed by almost all previous studies of autistic traits in adults.

To conclude, the present study provides a new perspective on the presentation of autistic traits in the adult population. Two distinct profiles were apparent following cluster analysis of a large general population sample, differing inversely in the magnitude of social and non-social (detail orientation) features of ASD. These differing profiles highlight a limitation in conceptualising ASD as a singular construct that varies across the population – and indicate, instead, that it is important to characterise adult individuals in their social- and detail-related traits independently. This coheres with the fractionable autism triad model better than a unitary spectrum view.

### Acknowledgements

This work was supported by an Australian Research Council Discovery Grant (DP1311336). J.H. is supported by an Australian Research Council Future Fellowship (FT100100322). P.E. is supported by a National Health and Medical Research Council Clinical Research Fellowship (546244). The authors declare that they have no conflict of interest.

## Linking text between chapters 5 and 6

In the thesis so far, predictive processing models of brain function have been discussed most directly in the context of *sensory perception*, *action*, and *statistical learning*. In addition, aspects of *social interaction* that are fundamental to the autistic phenotype were discussed briefly in *Chapter 1* with respect to sensorimotor formulations of predictive processing (i.e., active inference). In the social neuroscience of ASD research, a highly influential cognitive theory proposes that a specific deficit exists in *mentalising* or *theory of mind*; that is, the tendency to represent the *mental states* of other individuals, whether explicitly or implicitly (Baron-Cohen, 1997; Boucher, 2012; Frith, 2004; Happé, 1995; Pelphrey, Shultz, Hudac, & Vander Wyk, 2011). The mental states that others hold include their emotions, beliefs, and intentions, and our representation of these mental states are crucial for us to understand and anticipate others' behaviours, and thus for successful social interaction. Mentalising is a social-cognitive function of the brain that can be readily considered in the context of predictive processing. In particular, it is suggested that representations of others' mental states may arise via hierarchical unconscious inference on the external causes of sensory input (e.g., Kilner et al., 2007; Koster-Hale & Saxe, 2013; discussed in detail in *Chapter 6*). Linking hypotheses regarding aberrant predictive processing mechanisms to mentalising is thus an exciting avenue for understanding social and non-social aspects of ASD within the same neurocognitive framework, and for developing a more specific and neurocomputational understanding of the social-cognitive differences in ASD.

In *Appendix 2a*<sup>10</sup>, we discuss mentalising in predictive processing terms, and focus on the challenges to mentalising during particular forms of *social coordination* in which we need to establish *common knowledge* with other individuals. Common

---

<sup>10</sup> Published as Hohwy, J., Palmer, C. (2014). Social cognition as causal inference: implications for common knowledge and autism. In Mattia Gallotti and John Michael (Eds.), *Social Ontology and Social Cognition*, Springer Series "Studies in the Philosophy of Sociality", Vol. 4, 2014.

knowledge is a concept developed in the economic, game theoretic and philosophical literature, and, simply put, is the state in which two or more individuals have a shared understanding of their respective intentions and states of knowledge. Establishing common knowledge with others (e.g., via gaze signals) may be fundamental to coordinating our behaviour successfully with theirs, and for achieving the best individual outcomes in social contexts. We argue that atypicalities in how the brain infers the external causes of its sensory input in ASD may underlie mentalising difficulties, with particular implications for social behaviours that depend upon being able to establish common knowledge with others.

In *Chapter 6*, we focus more specifically on mentalising as an aspect of social interaction in the context of active inference.<sup>11</sup> A recent and exciting trend in social neuroscience research is to focus not only on the brain's response to social cues (e.g., participants' responses to gaze cues), but to examine cognition in situations that allow for reciprocal interaction between experimental participants (Schilbach, 2014, 2015; Schilbach et al., 2013). In typical social exchanges, our own actions directly modulate the cues that we receive from others, and thus a crucial aspect of social behaviour is how we actively investigate others (including their mental states) by interacting with them. In this regard, Schilbach et al. (2013) reviews evidence that neural networks engaged during social processing (for instance, in response to gaze cues; Schilbach et al., 2010) are notably distinct when we have the ability to reciprocally interact with others compared to when we merely observe them. Similarly, difficulties in social interaction are arguably central to psychiatric disorders, including ASD, and thus examining psychiatric differences specifically in the context of interaction may yield important insights into these conditions (Schilbach, 2016).

Thus, in phrasing social cognition in terms of implicit hierarchical inference, it is important to consider not only the passive inferential mechanisms that give rise to a representation of others' mental states, but also the inextricable role of action in

---

<sup>11</sup> See K. Friston and C. Frith (2015) and K. J. Friston and C. D. Frith (2015) for another treatment of social interaction in the framework of active inference.

the inferential process. In this regard, an important mechanism that has very recently been proposed as a supplement to predictive processing models of brain function is *counterfactual prediction*; that is, implicit predictions of the sensory consequences of a range of actions that we could perform in a given moment (but don't necessarily perform; outlined in *Chapter 1*). Counterfactual predictions are linked on a theoretical basis to action selection (Friston et al., 2012; Seth, 2015a) and perceptual phenomenology (Seth, 2014). Within a hierarchical model of unconscious inference on the external causes of sensory input, counterfactual predictions may be partly dependent on the inferred mental states of other people, and in this way play a crucial role in our social perception and in guiding our interactions with others. In *Chapter 6*, we make a link between counterfactual predictions and predictive processing models of social cognition, and consider the implications of this link for how we understand social interaction in predictive processing terms, the perceptual phenomenology of social experience, and the difficulties in social interaction in ASD.

*Chapter 6* was written for a special issue in *Consciousness & Cognition* edited by John Michael and Leon De Bruin that focusses on the *directness of social perception* (Michael & De Bruin, 2015).

## Declaration for Thesis Chapter 6

### Declaration by candidate

In the case of *Chapter 6*, the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
Developed ideas; reviewed literature; wrote the paper.	80%

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Anil Seth	Developed ideas; contributed to the writing.	N/A
Jakob Hohwy	Developed ideas; contributed to the writing.	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the candidate's and co-authors' contributions to this work.

Candidate's Signature		Date
		16/02/2016

Main Supervisor's Signature		Date
		16/02/2016

# Chapter 6

## The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism

Colin J. Palmer<sup>1</sup>, Anil K. Seth<sup>2</sup> and Jakob Hohwy<sup>1</sup>

<sup>1</sup>*Cognition and Philosophy Lab, Monash University, Clayton, VIC 3800, Australia*

<sup>2</sup>*Sackler Centre for Consciousness Science, School of Engineering and Informatics,  
University of Sussex, Brighton, BN1 9QJ, UK*

Published in *Consciousness & Cognition*, 2015

## Abstract

The mental states of other people are components of the external world that modulate the activity of our sensory epithelia. Recent probabilistic frameworks that cast perception as unconscious inference on the external causes of sensory input can thus be expanded to enfold the brain's representation of others' mental states. This paper examines this subject in the context of the debate concerning the extent to which we have perceptual awareness of other minds. In particular, we suggest that the notion of *perceptual presence* helps to refine this debate: are others' mental states experienced as veridical qualities of the perceptual world around us? This experiential aspect of social cognition may be central to conditions such as autism spectrum disorder (ASD), where representations of others' mental states seem to be selectively compromised. Importantly, recent work ties perceptual presence to the counterfactual predictions of hierarchical generative models that are suggested to perform unconscious inference in the brain. This enables a characterisation of mental state representations in terms of their associated counterfactual predictions, allowing a distinction between spontaneous and explicit forms of mentalising within the framework of predictive processing. This leads to a hypothesis that social cognition in ASD is characterised by a diminished set of counterfactual predictions and the reduced perceptual presence of others' mental states.

## 1. Introduction

There is an intuitive difficulty in how to characterise our experience of other people's mental states. At times we engage in reasoning about what might be going on in another person's mind, but it is also common for our awareness of others' mental states to appear more automatic and direct than explicit thought. For example, as our friend smiles at us we have an immediate sense of her happiness. As she glances at her drink we have an immediate understanding of her intention to reach out and grasp it. There is an ongoing contemporary debate regarding how 'direct' social cognition is (reviewed in Michael & De Bruin, 2015); the question is, to what extent do we have *perceptual* awareness of others' mental states? This subject is important, in part, because clinical conditions, including ASD and schizophrenia, are associated with atypicalities in social function that elude circumscribed descriptions in terms of perceptual and cognitive functions related to the representation of others' mental states. The neural and functional bases of symptoms in ASD (and schizophrenia) may be better understood as our conception of mentalising itself is further developed.

In this paper, we outline an account of how mental states come to be represented in the brain – this being mentalising – via the same process as perceptual representation. Namely, this is the unconscious process of hierarchical Bayesian inference on the (hidden) causes of sensory input. Just as non-social objects in our environment are causes of our visual input, the mental states of other people are a part of the physical structure of the world that produces the stream of sensory impressions that our brains receive. In this view, mentalising occurs implicitly and shares a fundamental similarity with the representation of non-social objects: each is a natural result of the brain's endeavour to best explain its sensory input.

Against this background, we argue that the perceptibility of mental states can be explicated in terms of whether they are experienced as having *perceptual presence*. This is the sense of veridicality attached to the content of our perception – the feeling that the objects that we perceive *really exist* in the environment around us. A recent development of unconscious inference points to specific functional

aspects of sensory processing that arguably underlie this sense of presence. On this view, perceptual presence rests on the hierarchical depth and – innovatively – the counterfactual richness of the generative models supporting unconscious inference, where counterfactual richness refers to the degree to which generative models encode predictions not only about the hidden causes of current sensory inputs, but also about the causes of prospective sensory inputs conditioned on possible but unexecuted actions. Generalizing these ideas to social perception furnishes a concept of ‘mental presence’, which rests on counterfactual predictions about the mental-state causes of sensory impressions. Consequently, we examine how difficulties in social cognition in ASD may relate to problems with the implicit predictive modelling of the sensory consequences of actions that we can perform to interact with others’ mental states.

In the next section, we cast mentalising as a process of causal inference and discuss how the representation of mental states can be situated relative to perception in the predictive processing framework of the brain. In sections 3 and 4, we show how the directness of social perception can be explicated in terms of the perceptual presence of mental states, and that this coheres with recent predictive processing accounts of this phenomenon that call on the counterfactual richness of hierarchical generative models. Finally, in section 5, we draw out the implications for our understanding of clinical conditions associated with differences in mentalising: is the perceptual presence of others’ mental states etiolated in ASD?

## **2. Social cognition as causal inference**

### **2.1 The world in the brain**

To understand the relationship between perception and social cognition, we should first consider how the brain comes to represent the world beyond the skull. To this end, we will briefly introduce the *predictive processing* framework of brain function (Clark, 2013; Friston, 2005; Hohwy, 2013). This framework provides a

powerful and compelling account of how the brain responds to sensory stimulation to embody information about its environment.

Predictive processing builds upon the notion of *unconscious perceptual inference* developed by Helmholtz (1860), Gregory (1980) and others, most recently in the form of Bayesian models of brain function (Kersten et al., 2004; Knill & Pouget, 2004; Vilares & Kording, 2011). This is the idea that what we are perceptually aware of, at any given moment, is the state of the world that is calculated as being *most likely* to be causing the sensory input that our brain receives, given prior beliefs about these causes that are furnished by previous experiences, development, and evolution. This idea emerges as a response to a fundamental problem, namely that through various causal interactions, physical states external to the brain (but including the body) conspire to produce an ambiguous stream of activity at our peripheral sensory receptors. Specifically, the relationship between states of the environment and sensory input is not one-to-one, as multiple causes interact non-linearly to produce the received input. The brain is thus faced with an inverse problem: the sensory effects of this process are directly accessible but their external causes are not. Yet it is only by controlling the causes of sensory input – which are physical features of the body and the world – that the brain can ensure adaptive behaviour and ultimately its own continued existence.

How then are the causes of stimulation represented in the brain? Predictive processing suggests that the problem is resolved by *hierarchical generative models* instantiated across the cerebral cortex (Friston, 2005). A generative model is a probabilistic mapping from worldly (and bodily) causes to sensory data; in other words, it specifies what sensory input *would* be received *if* a certain set of causes existed in the world. By feeding a generative model some expected (hypothetical) causes, we get *predictions* about the sensory data, which when mapped to cortical architecture are mediated by backwards connections between distinct levels in a hierarchy of cortical levels. These predictions are compared, continuously, to the *actual* sensory states at each level, with mismatch between the two constituting the *prediction error*. Prediction errors are propagated to higher cortical levels to provide information on how well the generative model and its input (the

hypothesised causes) are replicating the sensory data. *Recognition* is the process of determining what causes are most likely to produce the observed sensory input; this is unconscious perceptual inference, and occurs by implicitly ‘inverting’ the generative model by finding the hypothetical causes that minimise prediction error.

A crucial feature of this framework is that prediction-error minimisation occurs simultaneously at multiple levels of representation across the cortex (Friston, 2005; Shipp et al., 2013). Lower levels of the hierarchy encode hypothesised causes that act over smaller spatial and temporal scales (e.g., the perspective-dependent edges and colours that compose facial features). Higher levels represent expected causes operating over greater spatial and temporal scales (e.g., non-perspectival face recognition; Shipp et al., 2013). Expectations regarding higher-level sensory attributes generate predictions regarding the finer details represented at subordinate levels, thus constraining inference at lower levels with contextual information and generalised knowledge about regularities that occur in the sensory input. Expected causes encoded at higher levels are also revised to specify predictions that minimise prediction error. Thus, environmental (and somatic) objects are represented in a distributed and internally interactive manner.

The lowest level in this system is the sensory input itself – corresponding to retinal or lateral geniculate nucleus activity in the visual system, for example. The entire cortical hierarchy operating above shapes predictions arriving top-down at this level. This framework thus paints a compelling picture whereby the hierarchical causal structure of the world, in which causes occurring over shorter timescales interact with causes occurring over longer timescales to produce the sensory input, is recapitulated in the brain in the reverse direction. It is by minimising prediction error across the full hierarchy that the most likely expected causes of input are selected.

Representations of uncertainty and volatility are necessary for successful prediction and learning; for example, a change in sensory input could reflect either

a change in the *causes* of input or *variance* in the relationship between causes and their sensory effects. While there are multiple variables in this system that can be changed to minimise prediction error, the idea is that to fulfil this directive over the longer term, information about the true causal structure of the world must be approximated in the brain (Clark, 2013; Friston, 2005; Hohwy, 2013). It is important to appreciate that: (i) this is a process of inference that is *unconscious*, (ii) it is proposed to operate at the timescale of *perception* (e.g., as you shift your gaze to bring different objects into visual awareness, this is due to prediction error driving changes to the neural representation of the causes of visual input) and (iii) it entails ‘prediction’ of the sensory consequences of a suppositional model of worldly causes, rather than being limited to the identification of temporal sequences and the anticipation of future events.

Considerable further detail of this theory of cortical function is available in Friston (2005), including its consistency with a range of anatomical and functional facts and the computational theory underlying probabilistic inference. Similarly, the characterisation of the brain as an inferential system is discussed at length in Hohwy (2013) and Clark (2013). While the outline here has focussed on perceptual representation, the framework extends to encompass action (Adams, Shipp, et al., 2013; Shipp et al., 2013), attention (Feldman & Friston, 2010; Hohwy, 2012), interoception (Seth, 2013), and other cognitive functions: in fact, a strong motivation for considering this theory is its demonstrated ability to unify seemingly diverse brain functions within the notion of probabilistic inference on sensory states (Friston, 2010; Hohwy, 2013). We will now consider the advantages of understanding social cognition in these terms.

## **2.2 Mentalising as causal inference on sensory states**

A central ingredient of our social experience is that we represent the mental states of other people. In face-to-face interactions, for example, we commonly have an awareness of the beliefs, emotions and intentions of those around us. This sense of others’ mental states is a part of our understanding and anticipation of their behaviour, and moulds our own behaviour correspondingly. If our friend shows up

to the restaurant with a grim face, we have a sense of her mood and adjust our greeting accordingly. If she glances at our empty glass while pouring herself some wine, we have a sense of her intentions and might move our glass closer. This concept of *mentalising* is a dominant paradigm in the science of social cognition.

The mental states of other people are causes of our sensory input just as non-social objects are. Our friend's intention to pour the wine corresponds to a physical state in her brain that interacts with other causes in the world to modulate our sensory impressions. Her intentions might impact upon her eye gaze, speech and other bodily movements; her mood might cause changes in her facial expression, tone of voice and how she responds in conversation. These behaviours are in turn reflected in the sensory input that our brain receives. This is similar to the way our felt temperature is affected by a number of more or less deeply hidden causes throughout the day (nightfall, a sudden gust of wind, clouds, an open window, fever, putting on warmer clothes; and any interaction of such causes). In the predictive processing framework, inference regarding the most likely causes of input occurs as the brain tries to minimise prediction error across a hierarchy of temporal and spatial scales. Hence, we can conceive of mental states as just another feature of the world that comes to be modelled as the brain tries to best predict its sensory input. Just as we have a representation of the wine glass by virtue of its place in the causal pathway that determines retinal firing, we represent mental states because this too allows the brain to better predict its sensory input over time. Thus, mentalising slots into predictive processing as constituting the same kind of unconscious inference that the brain is already engaged in to represent its environment (Hohwy & Palmer, 2014; Kilner et al., 2007).

This is to say that mental state representations are expected causes of sensory input that have been probabilistically inferred. These causes are situated as part of a causal hierarchy, and share reciprocal interactions with higher and lower levels of representation. Thus, mental state inferences are statistically constrained by representations of longer-term expectations – perhaps regarding, for example, the kind of mental states that people tend to have in a given context, or the sense of

your friend's mood that has been reflected in a variety of her behaviours since she showed up to the restaurant, or even culturally defined social contexts and norms. Similarly, mental state representations are closely coupled to inference regarding lower level (and incontrovertibly perceptual) features of sensory input, with these hierarchical levels of representation continually adjusted to minimise prediction error throughout the system. By minimising prediction error, the most likely mental causes of the agent's actions can be inferred from observation, bypassing the inverse problem produced by the lack of simple relationships between observed behaviours and their mental causes (Kilner et al., 2007). For other Bayesian accounts of mentalising, see Kilner et al. (2007), Brown and Brune (2012), Zaki (2013), Diaconescu et al. (2014) and Koster-Hale and Saxe (2013), which include reviews of behavioural and neuroscientific literature supporting the utility of these frameworks.

### **2.3 What differs between mentalising and other forms of causal inference?**

Considering the view developed thus far, we should expect that what makes the representation of mental states different from non-social representation are the challenges involved in applying causal inference to this particular class of causes. In general, inference is made difficult by a lack of one-to-one relationships between causes and their effects. This is why we need probabilistic inference that incorporates prior knowledge about the likely states of the world and expectations regarding the degree of noise or uncertainty. There are several sources of ambiguity that are especially noteworthy in regard to mental causes. These are discussed briefly here and in more detail in Hohwy and Palmer (2014).

Firstly, mental states are relatively deeply hidden in the causal structure of the world. Mental states typically act on our senses via their effects on observed behaviour, and so have a less direct effect on sensory input compared to certain non-social objects. From a predictive processing perspective, mental state representations predict behavioural representations, which in turn predict lower level characteristics of the sensory input. This is to say that mental states are represented higher in the cortical hierarchy than these other features. Kilner et al.

(2007) sketch out a hierarchical representation of action observation along these lines, in the context of the mirror neuron system. In their framework, inferred intentions inform predictions about short-term motor goals, which in turn predict the motor commands that the observed agent is issuing, which in turn predict visual kinematics, which in turn predict lower level sensory features. By minimising prediction error across this hierarchy, a representation of the most likely mental state that underlies sensory input emerges alongside other levels of description, with each level of representation informed by both those higher and lower in the hierarchy via the reciprocal message passing that allows for prediction error minimisation.

Secondly, we can expect a significant role of context in determining the relationship between mental causes and their sensory effects. A colleague's glance at the clock during a meeting might reflect an eagerness to get home if it's late in the day, or an annoyance at my lateness in showing up if it's ten past nine in the morning. Deception and acting are more extreme examples of context dependency in the relationship between behaviours and the mental states that they reflect. It's worth noting again that this type of context dependency is not specific to social cognition. Consider the visual perception of a cat lurking in the bushes, its body partly obscured by branches, leaves and shade. At some level of representation, our brains infer that this is our cat that we are looking at, with, for example, an expected colour and visual form. However, the pattern of input that our brains should predict in the present moment is very much dependent on the immediate context of the cat's surroundings. In each case, causes in our environment operating over a variety of timescales interact non-linearly to produce our sensory input, and these interactions must similarly be modelled internally to most successfully match predictions to sensory data.

A third challenge to causal inference regarding mental states is more specific to the social domain. When we interpret other people's behaviour, we are often aware that they are also interpreting us and trying to model our own mental states. For example, in joint action, if you and I are trying to coordinate to move a heavy table into the other room, your sense of my intentions should depend on a

representation of *my* sense of *your* intentions. You know, for example, that I'm not going to try lifting the table until I think that you are also ready. The fact that other brains are modelling the world just as we are adds another degree of complexity to causal inference when applied to the internal states of other agents. This concept of 'meta-mentalising', discussed further in Hohwy and Palmer (2014), becomes particularly interesting in light of the need to establish *common* knowledge with other agents when we are trying to coordinate our actions with theirs.

In this section, we have conceived of mentalising as unconscious inference on sensory states; that is, a part of the process that also produces non-social perception. Moreover, this inference takes the form of prediction error minimisation, which is the matching between a hierarchical generative model of sensory states and the observed sensory states. This differs, of course, from the use of the term 'inference' to describe a higher-order, cognitive and consciously effortful process, which we also engage in at times when thinking about the mental states of other people. This points to a distinction between implicit and explicit forms of mentalising that will be developed further in Section 5. In the present section, we have also encountered specific challenges to inferring the existence and nature of mental states compared to that of non-social causes. These challenges help to define what is distinctive about unconscious inference of *social* states (or other minds), specifically. Casting mentalising in these terms is a useful step towards characterising how *perceptual* in nature social cognition is, which is the subject we now turn to.

### **3. The perception of mental states**

#### **3.1 Direct social perception in predictive processing**

A contemporary debate in cognitive science asks whether we are able to perceive the mental states of other people (reviewed in Michael & De Bruin, 2015). A starting point for this debate is the concern that mainstream accounts of social cognition (*theory theory* and *simulation theory*) assume that our awareness of

others' mental states are inferred from their behaviour in a manner that is distinct from perception (Bohl & Gangopadhyay, 2013; Gallagher, 2008). For instance, we may first have a perceptual experience of our friend's actions or facial expressions, and then, conditioned upon this, we come to an understanding of what kind of mental states underlie that behaviour. This extra cognitive step could involve a store of knowledge about the relationship between mental states and certain behaviours, or a kind of simulation of the mental states that observed behaviors would reflect were *you*, the observer, performing them. In disagreement with this picture, Shaun Gallagher presents a case for 'direct social perception': that our awareness of others' mental states is, in many instances, solely perceptual in nature (Gallagher, 2008; for related views, see, for example, Krueger, 2012; Zahavi, 2011).

To illustrate that perception can contain information that we might more intuitively think of as the product of higher-order cognitive processes, Gallagher describes his experience of non-social object perception:

*... when I open my eyes I see my car. It is true that it has a specific shape and is red, and I do see the shape and the color, but I see the shape and color as being aspects of something that is amazingly recognizable as my car. Actually, if you ask me what I see, I would likely not say that I see a red and shapely mass. Somehow I see through those aspects and I see my car. I do not see red mass, shape, and color, and then try to piece all of that together to make it add up to my car. I simply and directly see my car. (Gallagher, 2008, p.536)*

Accordingly, he argues, our perception of another's behaviour often *is* an awareness of that individual's mental state. As our friend smiles at us, we simply see that she is happy. As she reaches for the glass, we see her intention to grasp it. The thrust of this idea can be understood as defining our sense of others' mental states as a form of recognition associated with the perceptual content of observed behaviour, or a higher-order level of description of this content. Gallagher's position includes two closely related claims: (i) that perception can contain an awareness of others' mental states and (ii) that the involvement of a cognitive

process going beyond perception (e.g., interpretation or explicit inference) isn't necessary for us to be aware of another's mental state. While the basis for these claims is largely phenomenological, empirical research centered on the sensitivity of infants to emotional and intentional social stimuli is also taken to suggest automaticity to social experience or an independence from developed cognitive abilities (Gallagher, 2008).

Gallagher argues against the notion of there being hidden mental causes, and so in this respect the predictive processing view on mentalising we develop is very different from Gallagher's in that we still conceive of mental states as hidden causes of sensory data. However, a predictive processing account of direct social perception is still conceivable and consistent with much of Gallagher's phenomenological description of social perception. This is because, in the Bayesian view of perception, processes of unconscious inference can and do give rise to 'direct' perceptual experiences of inferred hidden causes. In other words, non-social perceptual objects are inferred from sensory data in just the same way that mental states are – any differences between how these states of the world are represented relate to the challenges involved in causal inference on different kinds of causes.

Predictive processing thus provides a different picture of inference than that which Gallagher discerns in theory-theory and simulation theory. Specifically, the extra-neuronal causes of sensory activity are represented in a hierarchical manner, with higher levels more divorced from the present sensory input, reflecting their association with causes operating over larger spatiotemporal scales. The relationship between mental states and observed behaviour is thus comparable to the relationship between local and global perceptual features, or causes that modulate input over shorter and longer timescales. Moreover, this relationship is bidirectional. The inference of a particular mental state entails predictions about lower level representations, potentially shaping lower level perception; conversely, prediction errors update mental state representations to better account for lower-level features. Thus, lower-level perceptual features are not merely the evidence that inference is conditioned on, in order to produce a mental

state representation, but rather, mental state representations and lower-level representations are each adjusted in interaction with one another to minimise prediction error over time. Gallagher (2008) argues that, "The question about direct perception, whether of objects or of others is ... [about] how smart, how richly informed, it is. The smarter the perception is, the more work it does; the dumber it is, the more it requires extra cognitive processes (theory, simulation) to get the job done." The view of mentalising as a level of representation in a predictive cortical hierarchy provides us with a functional architecture of unconscious sensory processing that could underlie the 'smart' perception of mental states; that is, inferred representations of behaviour that are informed by, and closely tied to, the mental states inferred at levels above. This helps to specify the relationship between lower-level perceptual representation and mental state representations, while circumventing the phenomenological criticism of inferential accounts that mentalising doesn't require any interpretative cognitive processes that go beyond perception.

Having said this, casting mentalising in predictive processing terms does not necessarily entail that mental state representations are experienced in a perceptual manner. While predictive processing describes how information that we are perceptually aware of comes to be encoded in the brain, it isn't necessarily the case that we have the same kind of perceptual awareness for all of the modeled causes of input represented across the cortical hierarchy. In particular, there is a spectrum between shorter- and longer-term expected causes of the sensory input, grading between, for example, more perspectival and more perspective-invariant levels of representation. Experientially, different levels of representation could be more or less perceptual. Similarly, while Gallagher provides an argument that observed behaviours can be automatically recognised as intentional or emotional states, the lack of a firm distinction between perceptual experience and the experience of (automatic) knowledge associations, for example, means that ambiguity still remains in how we experience the mental states of others.

### 3.2 Do mental states have perceptual presence?

In the previous section, our experience of others' mental states was depicted (in our reading of Gallagher) as a type of recognition that is closely entwined, phenomenologically, with perceived behaviour. This helps to clarify the perceptibility of others' mental states, while also answering in part why there is some ambiguity in this regard – as our experience of mental states is defined as a type of 'higher order' perceptual property as distinct from the rawer perceptual content also associated with the observed behaviour (like the "red mass, shape, and color" of the car). There is more to perceptual experience, however (and to the intuitive question of how perceptual our sense of others' mental states is) than how automatically and effortlessly our awareness comes about.

In particular, we can ask whether the recognition of a mental state in observed behaviour is experienced as knowledge 'in our head' or as a veridical state of the external world. Hence, a crucial quality of perception to consider in this debate is the *perceptual presence* or *subjective veridicality* of others' mental states:

*In normal circumstances perceptual content is characterized by subjective veridicality; that is, the objects of perception are experienced as real, as belonging to the world. When we perceive the tomato we perceive it as an externally existing object with a back and sides, not simply as a specific view—a "perspectival take"—on an external scene. (Seth, 2014, p.2)*

When we look around us, at the table, at the coffee cup, or at another person, we have the experience that these are real objects in our environment. Subjective veridicality, "whether the perceptual content appears, phenomenologically, as part of the external world" (Seth, 2014), can be contrasted to *doxastic veridicality*, for which "perceptual content is understood cognitively to reflect part of the external world", and further to *perceptual reality*, "vivid perceptual phenomenology" that

doesn't necessarily include an associated sense of presence, but which rather marks a difference between perception and imagery or imagination.<sup>12</sup>

Interestingly, certain perceptual experiences seem to lack perceptual presence, thereby motivating a principled account of why some experiences would be associated with a sense of presence and others not. Synesthetic concurrents and certain kinds of hallucinations are examples in this respect (discussed in Seth, 2014), but an illustration that is typically more readily available is that of visual afterimages. If you gaze at a bright light and then look away you will experience a kind of residual watermark in your visual field. There is certainly perceptual content – you can really see the afterimage in your field of view – but it isn't perceived as a robust part of the external world; one could say that afterimages have diminished or absent perceptual presence.

Do we experience the mental states of other people as belonging to the world? That is, do mental states have subjective veridicality or only doxastic veridicality? While this issue isn't explicitly discussed in Gallagher (2008), the direct social perception account suggests that it is introspectively plausible that our sense of others' mental states has subjective veridicality in some instances. Moreover, as described in the following section, presence should be expected for our awareness of mental states when these phenomena are considered within the predictive processing framework. Subjective veridicality may be a crucial concept for defining how perceptual mental state representations are, and why in some instances mentalising may be experienced more perceptually while in other cases as a

---

<sup>12</sup> We thank an anonymous reviewer for pointing out that discussion of the subjective veridicality of perceptual entities can be traced back through earlier psychological literature to Albert Michotte's concept of 'phenomenal realness' (Mausfeld, 2013; Thinès, Costall, & Butterworth, 1991), "the impression that a perceptual entity is perceived as having an autonomous existence in our external mind-independent world" (Mausfeld, 2013). Michotte richly describes the experience of certain visual illusions, flat images of three-dimensional objects (drawings, photographs, etc.) and other phenomena to illustrate that phenomenal realness is a separable attribute of our perceptual experience that can be experienced more or less strongly for different perceptual entities. The development of these ideas in terms of sensorimotor contingencies and the counterfactual predictions of hierarchical generative models (Seth, 2014), described in Section 4.1, provides a mechanistic explanation of this attribute.

process of explicit cognition. This approach seems to capture a part of Gallagher's phenomenological intuition that mental states are experienced without any 'additional interpretation or inference': that is, mental states with perceptual presence are sensed as being states of the world around us, not as mere ideas or interpretations in our head.

An initial objection to the idea that mental presence has much to do with the directness of social perception is that presence and perceptual reality can come apart. For instance, it could be argued that we don't need any cognitive process beyond perception to have awareness (perceptual reality) of visual afterimages, despite their relative lack of perceptual presence, and so should conclude that presence is not necessary for direct perception. However, as discussed, Gallagher's notion of social perception relates to what we have called higher order perceptual properties (such as recognition or perceptual integration), rather than perceptual reality (like the colour and shape of an afterimage; see also Bohl & Gangopadhyay, 2013 for a related point about mental states). Thus, our focus on presence is built upon the view that the intuitive question of 'how perceptual' mentalising is, relates in part to how we characterise the experience of these higher order properties of perception; whether we experience mental state representations as being subjectively veridical is a crucial aspect of this matter.

Similarly to Gallagher, the authors Scholl and Gao (2013) contrast perceptual processes with conscious (explicit) inference or decision making, while acknowledging the difficulties in drawing a firm line between perception and cognition. Importantly, they argue that "a hallmark feature of perception (vs. cognition) is its strict dependence on subtle visual display details; percepts seem to be irresistibly controlled by the nuances of the visual input regardless of our knowledge, intentions, or decisions." Scholl and Gao (2013) review phenomenal, psychophysical and neuroimaging evidence regarding our detection of social agents in visual stimuli, famously illustrated, for example, in the sense of *animacy* or intentionality that can be elicited even by simple geometric shapes when they are moving appropriately (e.g., "the triangle is fighting the square"; Heider & Simmel, 1944). These authors ask whether our appreciation of animacy in visual

stimuli is perceptual in the sense of reflecting a property of visual experience comparable to colour, shape, and orientation, or instead, whether attributions of animacy stem from explicit decisions conditioned upon perceptual experience. In their experimental work, they exploit an implicit measure of agency detection: the ability of a participant to avoid a geometric shape that is ‘chasing’ a participant-controlled avatar across a computer screen (against a background of similar non-chasing objects). A systematic relationship was found between the variability of the antagonist’s motion and the participants’ difficulty in detecting its chasing behaviour. The subtlety of the stimulus features that modulated agency detection in this task was taken as evidence that our awareness of animacy is not dependent on conscious appraisal of the stimuli, but is instead a perceptual process.

The view of Scholl and Gao is consistent with the predictive processing account of social cognition, for which awareness of social states of the environment arises through the same mechanisms as the perception of non-social environmental properties (i.e., unconscious inference). Our focus on perceptual presence in this paper (rather than on perception *per se*) builds on the notion developed in Seth (2014) and Thinès et al. (1991), for example, that presence is a separable property that perceptual entities can be associated with to a greater or lesser degree. Further, while we have argued that the predictive processing account of social cognition is fully consistent with direct social perception, the former account doesn’t necessitate that (unconsciously) inferred social states are experienced as perceptually as are visual properties like shape and orientation. Rather, we emphasise that the same unconscious mechanisms related to inference are crucial to non-social perception and mentalising, including those that may underlie perceptual presence, detailed in the next section.

## 4. Perceptual presence and predictive processing

### 4.1 Perceptual presence in the brain

Predictive processing is in part a theory of how neurocognitive processes account for perceptual content, its phenomenological properties and interactions with other cognitive processes like attention and memory. To help account for the phenomenon of perceptual presence, Seth (2014) expands upon this framework to incorporate *counterfactual predictions*. The suggestion, drawing on Friston et al. (2012), is that the hierarchical generative models implemented in the cortex include predictions about how the sensory input *would change were* we to interact with the world in the various possible ways that we can. For a given point in time, this is the case for both the actions that we *will* actually perform next and, importantly, a range of alternate actions that we *could* perform but won't in this instance. To illustrate, when we look at the coffee cup on our desk, we receive retinal input from one side of the cup, and under predictive processing, the brain hypothesises a worldly cause (or hierarchical set of causes) that accounts for this input. The further suggestion, however, is that the brain is engaged in predicting the sensory consequences of interacting with this object: like moving around the cup in various ways, or picking up the cup, or occluding our view of the cup with another object. This development of predictive processing provides a neurocognitive operationalization of the notion of 'sensorimotor contingencies', the mastery of which has previously been associated with perceptual presence (O'Regan & Noë, 2001). The argument detailed in Seth (2014) is that examples of perception that involve reduced perceptual presence (e.g., visual afterimages, synesthesia concurrents and certain hallucinations) are likely to be associated with an impoverished set of counterfactual predictions, as there are fewer ways in which we can act to change our input contingent on these kinds of (hypothesised) worldly causes. Correspondingly, perceptual presence comes about by virtue of there being a *rich* repertoire of counterfactual predictions associated with the (hierarchical) representation of the object in question. There are many ways in which we can interact with the coffee cup to change our input – this accounts for why our perception of the cup is associated with a strong sense of presence. There

are fewer ways in which we can act to change our input based upon the visual afterimage – thus the brain's representation of the afterimage is associated with fewer counterfactual predictions and, experientially, the afterimage has reduced presence.

The counterfactual richness hypothesis of perceptual presence is built in part on the unusual and revealing phenomenology of synesthesia. In synesthesia, the perception of an *inducer* stimulus elicits a *concurrent* percept, either of the same or another sensory modality. Significantly, while both inducer and concurrent percepts are experienced with perceptual reality, only the inducer is experienced as being subjectively veridical (reviewed in Seth, 2014). The counterfactual richness hypothesis assigns a separable mechanism of hierarchical generative models to perceptual reality (i.e., the inferred causes of sensory input and the corresponding non-counterfactual predictions) and perceptual presence (i.e., the richness of hierarchical counterfactual predictions). This furnishes a neat account of why these two aspects of perceptual content tend to co-exist but come apart in the case of synesthesia concurrents. Specifically, it is precisely because concurrents reflect hypothesised states of the world that synesthetes are unable to interact with, that they are resistant to the formation of a rich hierarchical set of counterfactual predictions. In contrast, while, like concurrents, the perceptual reality of inducer stimuli reflects their status as an inferred cause of sensory input, the sensory consequences of our actions are contingent upon inducer stimuli to a much greater degree. Thus, the brain is able learn these sensorimotor contingencies over time and, in a given moment, model a richer set of counterfactual predictions for inducer stimuli. The hypothesis that perceptual presence relates specifically to the action-dependent prediction of sensory outcomes similarly allows us to account for the characteristics of a range of other types of perceptual experience, such as imagery and hallucinations (discussed in Seth, 2014).

The notion of counterfactual prediction of sensory states builds upon a more elementary extension of the predictive processing framework to action, known as 'active inference'. The goal of the brain under this framework is to minimise

prediction error: this can be achieved both by changing predictions to match the observed data and, via action, changing the sensory input to match predictions. Thus, movement is cast as a process by which initially *inaccurate* predictions regarding the parameters of proprioceptive feedback (mediated by descending cortical projections) are fulfilled by spinal reflexes that modulate muscle activity to bring proprioceptive states into line with these predictions (Adams, Shipp, et al., 2013; Shipp et al., 2013). In this manner, prediction error minimisation is an ongoing synthesis of active and perceptual inference. As well as specifying how action occurs, the concept of active inference hangs together with the idea developed by Helmholtz (and later, Richard Gregory) that unconscious inference is an investigative process, where intervening on the sensory stream allows us to test hypotheses regarding the state of the world (Friston et al., 2012; Gregory, 1980; Helmholtz, 1860). Predicting the sensory consequences of a range of *possible* actions (i.e., counterfactual prediction) is then important to allow action selection in a manner that will optimise our inferred representation of the world by reducing uncertainty (Friston et al., 2012; Seth, 2014, 2015a).

The expected precision of predictions plays a key role in both action selection and perceptual inference. It is those counterfactual predictions that are most highly precise that will initiate action via active inference. In the case of perceptual inference, the expected precision of (non-counterfactual) predictions and prediction errors also regulates how perceptual content changes: in particular, prediction errors are more likely to cause updating to representations of the world if the sensory evidence is expected to be highly precise relative to existing predictions. In this way, the counterfactual richness hypothesis of perceptual presence differs in its emphasis on precision expectations in determining perceptual content, as counterfactual predictions that are less precise (such that they do not end up driving action) still contribute to phenomenology in the form of perceptual presence. Moreover, the *balance* between active and perceptual inference is regulated by the relative precision of predictions. Thus, it is conceivable that a rich set of counterfactual predictions (and hence perceptual presence) can exist when simply viewing an object without acting upon it, despite the intimate relationship between action and counterfactual prediction. See

Hohwy (2014) and Seth (2015b) for further discussion regarding the status of counterfactual predictions.

An alternative explanation for perceptual presence in predictive processing terms focuses instead on the role of higher-order perspective invariant object representations (Hohwy, 2013, 2014). In this view, the sense of presence is related to the representation of external objects as persisting in time, encoded at superior levels of the distributed hierarchy of representation than perspectival takes on the object<sup>13</sup>. This account also relates in part to the role of action in predictive processing: perspective invariant representations are presumably built up as we experience an object from multiple perspectives, or, in other words, as persisting to some degree as an external cause of stimulation in the face of the interventions that we can perform upon the sensory stream. The functional and neural instantiation of this account differs from the counterfactual view, however; presence relates here to ‘passive’ components of representation that exist higher in the cortical hierarchy than perspectival representations, while for the view described in Seth (2014) presence relates to action-dependent counterfactual predictions that are distributed across the cortical hierarchy, including at perspectival levels of representation.

These two accounts are not mutually exclusive, and might be reconciled. For example, higher-order invariants may support counterfactual predictions (at corresponding and lower levels), and, further, developing counterfactual predictions at lower (perspectival) levels may support the development of higher-order invariants. See Seth (2015b) for further discussion of this point. The distinction between these two accounts is in part between perceptual inference and active inference, processes that are closely intertwined in the effort towards

---

<sup>13</sup> An anonymous reviewer questions whether invariance can help to explain presence given that we seem able to think of things invariantly without an associated sense of presence. We stress that under this account presence occurs when large swathes of the hierarchy are activated simultaneously: that is, precisely when variant and invariant levels of the hierarchy together can suppress prediction error. Thought without presence may be when predominantly higher levels are activated (and perception without presence is when predominantly lower levels are activated).

improving the brain's model of the world via prediction error minimisation. Similarly, the sense of presence might be built up both via passive observation (e.g., watching two colleagues talk) and active observation (joining in on the discussion to probe the mental states that have been passively inferred).

## 4.2 Mental states and counterfactual predictions

We can hypothesise that counterfactual predictions occur across multiple levels of the predictive processing hierarchy: for example, perspective invariant representations of a face may allow counterfactual predictions regarding global features of the face, which in turn allow counterfactual predictions regarding local details of the face. In fact, counterfactual predictions should be expected to exist in relation to any level of the external causal hierarchy that our actions can interact with to modulate our sensory input in a discernable manner, thus enabling active inference. Thus, an implication of the counterfactual account is that subjective veridicality can potentially be experienced for any level of causal representation, depending on the extent of associated counterfactuals.

We have already seen how mentalising can be conceived of as a level of representation of the external causes of sensory input in a hierarchy of predictive processing. We can now ask whether mental state representations are likely to be associated with counterfactual predictions. This immediately seems plausible for many instances of mentalising. If our friend is unhappy, this allows predictions about how the sensory input would change if I were to make a joke. If she is intending to reach out for the wine bottle to fill up her glass, this should shape predictions about the sensory effects of moving the bottle away from her, or moving my glass closer. Face-to-face conversation in particular illustrates a rich sequence of auditory and visual input that is quite directly modulated by our own utterances and expressive behaviours, the immediate sensory consequences of which are highly dependent on the succession of beliefs, intentions and emotions that we become aware of in our conversational partner throughout. Conversation is thus a process where we actively interrogate the mental causes of our sensory input, actively sampling new data by modulating the facial expressions of others,

their bodily actions, and their speech. The immediate sensory consequences of our actions – things we say, where we look, movements we make – are predictable in part based on the inferred mental states of those around us, and thus it is conceivable that mental state representations are associated with a broad repertoire of counterfactual predictions. This is to say that, as we observe our friend, the sensory consequences of a range of various possible actions that we could perform are predicted dependent in part on the mental states that we model her as having.

This contrasts with other instances of mentalising for which we are less able to intervene upon the world to change our sensory input in a way can be predicted dependent upon the mental states of other people. Interestingly, in this regard, the temporal scale of representation is one factor that may determine the extent of the accompanying counterfactual predictions. For example, if I receive an email containing the schedule of exam times for this semester, I might experience the notion that my friends will be stressed out in a months time when the exam period starts, but due to the timescale involved this may not be associated with a rich set of neural predictions regarding how my input will change consequent upon actions I can currently perform. Action necessarily occurs over certain shorter timescales; thus, the representation of causes that operate over longer time scales may tend to have diminished subjective veridicality than those that allow for counterfactual predictions over the timescales that our actions can more directly impinge upon.<sup>14</sup>

---

<sup>14</sup> The physical constraints of our body determine the repertoire of possible actions that we can perform, and the types of causes in the world that we are able to interact with. Our bodies therefore restrict the counterfactual component of our brain's representation of the external world. For example, we perceive the clouds in the sky, but we are limited in how we can act to change our sensory input consequent upon them, which, if the counterfactual view of presence is correct, determines the extent of veridicality with which they are perceived. This is to say that the quality of our perceptual experience of the world is dictated not only by the limitations of the sensory organs that determine what the brain can represent about the physical world (e.g., their stimulus selectivity, resolution, location, etc.), but also by the limitations of our actions in intervening upon the causes of our sensory input. There are interesting issues here regarding relations between perceived presence and perceived objecthood; see Seth (2015b) for more discussion.

Similarly, certain kinds of stimuli that elicit mental state attributions can't be interacted with as fruitfully as the stimulation that is received in face-to-face interactions. For example, as I read an email, or view a schematic version of the Sally-Anne task, or try to decipher a New Yorker cartoon, there are relatively few ways in which the mental states I attribute (to the absent author of the email, the fictional characters in the cartoons, or perhaps the cartoonists) allow predictions about how the sensory input will change consequent upon the actions that I can currently perform. This contrasts with dynamic stimuli such as the gestures and facial expressions of a person that we are directly interacting with face-to-face: here my actions very much determine my consequent sensory input in a way that is dependent on the mental states that modulate these stimuli.

The counterfactual account of presence implies that "perceptual presence will be lacking when the corresponding generative models are counterfactually poor" (Seth, 2014). This provides a principled way of characterizing the degree to which mental state representations are associated with this dimension of perceptual phenomenology. Representations with a richer set of counterfactuals, whether due to the timescale or type of stimuli involved, are more likely to be experienced as veridical states of the external world than those with a poorer set of counterfactuals. This is important, because it allows us to hypothesise a distinction between different types of mentalising in terms of the corresponding counterfactuals (and sense of presence), discussed further in Section 5 with respect to implicit and explicit forms of mentalising. It may be that mental states associated with a poorer set of counterfactuals are experienced as conceptual associations or explicit knowledge, while those with a richer set (e.g., those that occur during conversation) are experienced as subjectively veridical states of the surrounding perceptual world. This point might generalize such that the difference in our experience of certain explicit thoughts or knowledge from perception (e.g., our abstract sense of the Andromeda galaxy, or our knowledge that the earth is spherical) can be explained in part in terms of a concomitant weak set of counterfactual predictions.

### 4.3 Alternative theories of perceptual presence and mentalising

In the discussion of perceptual presence in the literature, one focus is on how hidden parts of visual objects (e.g., the side of the coffee cup that is presently facing away from you) contribute to our perceptual experience of these objects (Noë, 2004). That is, as we look at the cup only the near side of it impacts upon our sensory input, but nevertheless we are perceptually aware of the cup as a three-dimensional object with components that are currently unseen. An analogy in this regard has been made previously by other authors to our perceptual experience of mental states (Bohl & Gangopadhyay, 2013; Krueger, 2012; Smith, 2010): drawing upon Husserl's characterisation of 'present' and 'co-present' aspects of perceived objects, this earlier idea is that we may experience mental states along with our experience of observed behaviour in a manner similar to how the hidden parts of the cup contribute to our experience of the cup as a 'full' object even as we have a more direct experience of the near side. It might be possible to develop this view in line with the notion of counterfactuals in hierarchical generative models: however, there is an issue here in how mental states can be modeled as counterfactual causes of the sensory input (analogous to the hidden sides of the cup) without also being modeled at times as the non-counterfactual causes of input (i.e., the causes of the present input). Indeed, this kind of account is inspired partly by the notion that mental states are hidden such that they can't be as directly verified as non-social objects (Bohl & Gangopadhyay, 2013).

The view developed in the present paper differs importantly from these previous accounts in that mental states are represented as distinct causes of our sensory input that are separate from, and that modulate, more proximal causes such as another's behaviour. That is, mental states are more closely analogous to the *near* side of the cup in that they can be modeled as (non-counterfactual) causes of the present sensory input – this conceptualization stems from our view of worldly representation, whether concerning mental states or visual objects, as being a process of causal inference on sensory states within a hierarchical generative model. Further, we suggest that mental state representations themselves are in some instances accompanied by a sense of subjective veridicality (and

counterfactual predictions). The predictive processing view avoids some potential criticisms of the idea that the brain models mental states as direct causes of sensory input (see Bohl & Gangopadhyay, 2013). Firstly, this view avoids trending towards behaviourism, because mental states are still represented distinctly from observed behaviour. Secondly, it provides a principled way of accounting for the intuition that mental states are more hidden than non-social objects like the cup, in terms of the challenges to performing causal inference on mental causes discussed in Section 2.3.

## **5. Autism and counterfactual predictions regarding others' mental states**

### **5.1 Implicit and explicit mentalising in ASD**

Frith and Frith (2008) contrasts social cognition that entails “fast, relatively inflexible routines that are largely automatic and implicit and may occur without awareness” with that involving “slow, flexible routines that are explicit and require the expenditure of mental effort.” One example of this division between implicit and explicit processes occurs in the developmental trajectory of mentalising. A cornerstone of research in this area concerns laboratory tests for the ability to represent the *false beliefs* of others; for example, a participant may be shown an actor, Sally, hiding Anne’s ball while Anne is away in the other room – to understand where Anne will search for the ball when she returns, the participant must have some sensitivity to how Anne’s state of knowledge differs from their own. There is evidence from developmental studies that children demonstrate a kind of implicit knowledge of others’ false beliefs (reflected in gaze direction and looking time) even at an age when their verbal predictions of the actors’ behaviour fail to demonstrate an explicit knowledge in this regard (Clements & Perner, 1994; Frith & Frith, 2008). Comparable forms of looking behaviour are also apparent in (preverbal) infants (Onishi & Baillargeon, 2005; Surian, Caldi, & Sperber, 2007).

This distinction between implicit and explicit forms of mentalising has played a crucial role in clarifying the deficits in social cognition that occur in ASD. This is a diagnosis that covers a diverse set of individuals who show significant differences from one another in their social cognition, language and intellectual functioning (Lai et al., 2014; Lai et al., 2013). Frith (2004) suggests that a fundamental deficit that exists across this spectrum is in the implicit or spontaneous sensitivity to the mental states of other people, with socially higher-functioning individuals more able to compensate for this deficit by drawing on explicit reasoning and learning skills. Support for this idea comes, for example, from the examination of eye movements in a group of adults with high-functioning ASD while they completed a false belief task (Senju et al., 2009). These subjects were inseparable from age-matched controls in their responses to a variety of tests of mentalising ability that were conveyed to the participant verbally – one example of which was a version of the Sally-Anne scenario sketched above. Nevertheless, in a *nonverbal* false belief task that was acted out with puppets by an experimenter, the behaviour of the groups diverged. Specifically, controls were significantly biased in their direction of gaze towards stimuli made salient by the implied false beliefs of the actors, while those with ASD were not, suggesting that those with a diagnosis lack a spontaneous sensitivity to the mental states of other people, becoming most apparent when prompts to explicit cognition are unavailable.

Tager-Flusberg (2001) makes a comparable proposal to Frith (2004) regarding social cognition in ASD, contrasting ‘social-perceptual’ and ‘social-cognitive’ aspects of theory of mind. Gallagher similarly suggests that ASD may deviate from typical social cognition in that mental states are not as readily experienced ‘directly’ (i.e., as given in our perceptual experience of others’ behaviour), but rather that awareness of others’ mental states tends to require the (explicit) interpretation or reasoning that, he argues, those without ASD more commonly limit to social situations that are unusually ambiguous (Gallagher, 2004, 2008).

Importantly, this notion of a more specific deficit in ASD in implicit, spontaneous or perceptual forms of mentalising can be developed in light of the idea of perceptual presence and counterfactual predictions in predictive processing.

Specifically, we can hypothesise that *mentalising in ASD fails to occur spontaneously as a part of the predictive inference on sensory states that operates on the timescale of perception, and instead relies upon a slower process of explicit reasoning*. Importantly, this may entail that the representation of mental states in this condition is less directly associated with the unconscious modelling of counterfactual predictions.

Firstly, this implies that others' mental states are experienced with reduced perceptual presence in ASD compared to those without this diagnosis: that is, the thoughts or feelings of others are more likely to be experienced as conceptual associations or explicit interpretations rather than as veridical properties of the surrounding world. Section 4 described how, in general, the quality of subjective veridicality is expected for mental state representations that are replete with counterfactual predictions, like, as we suggest, our sense of others' mental states in face-to-face conversation. Thus, it is these social situations that would be experienced most differently in ASD, rather than situations that, even in those without social cognition deficits, would tend not to elicit a rich set of counterfactual predictions due to the timescale or type of stimuli involved. As discussed in Section 4.2, the latter is likely to include mentalising based on written stimuli (like stories, or email correspondence) as opposed to more dynamic social interactions. This coheres with findings that higher-functioning individuals with ASD (particularly adults) respond similarly to controls in traditional written or verbal tests of mentalising ability (e.g., Callenmark, Kjellin, Ronnqvist, & Bolte, 2013; Senju et al., 2009), despite experiencing social difficulties to the extent necessary to meet a diagnosis of ASD.

Moreover, counterfactual predictions are a basis on which action selection occurs under predictive processing, in the setting of active inference. This provides a link from the idea of reduced spontaneous mentalising in ASD to consequences for action and behaviour. In typical function, the implicit representation of mental states may come with the unconscious modelling of counterfactual outcomes, important for driving our behaviour in a manner that is sensitive to the mental states of other people. Specifically, on the basis of counterfactual predictions,

actions are selected to minimise uncertainty in our inferred representation of the world (Friston et al., 2012; Seth, 2014, 2015a). As we described earlier, the mental state causes of our sensory input are actively interrogated via our own verbal and expressive behaviours. Thus, if the brain were not unconsciously modelling the counterfactual outcomes of our possible interactions with others' mental states, this may incur a profound impediment to the selection of actions that optimise the inferred representation of the world by reducing uncertainty (contingent on the mental states of others). Actions that are sensitive to the sensory consequences of interacting with others' mental states would not be generated as spontaneously via predictive processing, and perhaps instead require top-down, deliberative control. This is to say that a deficit in implicit mentalising in ASD, when this is cast under the predictive processing umbrella, is not just about the perceptual awareness of others' mental states, but also the unconscious processing that underwrites the selection of actions that fine tune generative models. A lack of counterfactual modelling that would otherwise drive mentalising-sensitive action may underlie the disassociation between eye tracking and verbal response results of Senju et al. (2009), described above.

This develops the emerging picture that adults with ASD can 'do' mentalising, but lack important automatic aspects of the process as compared to non-ASD people, where these aspects rest on the counterfactual predictions in hierarchical generative models. Thus, drawing on explicit reasoning skills may compensate in part for deficits in spontaneous mentalising, but fail to ameliorate social difficulties due to an unresolved absence of implicit counterfactual predictions.

## **5.2 Predictive processing in ASD**

The inherent ambiguity between sensory input and its causes (necessitating inference) allows for the possibility of different models of the world arising in different individuals, contingent upon our different sensory streams. The inexhaustible richness of worldly causes and their interactions, some subset of which is modelled in the brain, similarly helps to explain why individual differences in accounting for the sensory data would arise via predictive

processing. Several hypotheses have been recently put forth regarding how causal inference may occur differently in ASD, contributing to social and non-social symptoms alike (Hohwy, 2013; Hohwy & Palmer, 2014; Lawson et al., 2014; Palmer et al., 2013; Palmer, Paton, Kirkovski, et al., 2015; Pellicano & Burr, 2012; Skewes et al., 2014; van Boxtel & Lu, 2013b; Van de Cruys et al., 2014). These hypotheses help to explain why the unconscious inference of others' mental states may be compromised in ASD, resulting in a reliance on explicit cognition and reduced counterfactual modeling in social situations.

### 5.2.1 Precision expectations in ASD

A core mechanism of predictive processing that has been drawn upon to understand ASD is the *weighting* of prediction errors, via top-down modulation of sensory processing. This weighting determines the extent to which new sensory data drives changes to the brain's representation of the world. Such a mechanism is necessary because changes in sensory input could reflect either real changes in the underlying causes of input, or various sources of context-dependent variability or noise in the mapping from a given set of worldly causes to their current sensory effects. To arrive at an accurate model of the external world, the reliability or precision of sensory estimates (as the inverse of variance) must be estimated. Hence, the weighting of prediction errors is suggested to occur in proportion to the *state-dependent expected precision* of sensory information relative to the expected precision of existing predictions (Feldman & Friston, 2010). Determining how much precision to expect in sensory states across different contexts is itself a matter of learning and inference suggested to come about via prediction error minimization. Expecting low precision in the sensory input carries the risk of persisting with a given model of worldly causes in the face of fluctuations in the sensory input that reflect real changes in the environment (treating these fluctuations as noise instead). Expecting high precision carries the risk of modeling the expected causes of input on uninformative fluctuations in the input (noise) and failing to generalize across samples to arrive at a deeper and more predictive model of the underlying causes of sensory data (i.e., overfitting the model to sensory data; Hohwy, 2013).

Specifically, in relation to ASD, a common hypothesis that has emerged is that sensory processing is associated with an expectation for high precision in sensory input, related to the processing of state-dependent levels of uncertainty (Hohwy, 2013; Lawson et al., 2014; Palmer et al., 2013; Palmer, Paton, Kirkovski, et al., 2015; Van de Cruys et al., 2014). Physiologically, this is thought to be constituted in the form of greater gain on cortical neurons signaling prediction error, where this gain is controlled context-dependently by top-down modulation mediated by several interacting neuromodulatory systems (Lawson et al., 2014; Quattroki & Friston, 2014). Importantly, reduced sensitivity to the ambiguity inherent in inference is suggested to instill a reliance on more detailed, less generalized, or lower level representations of the world, as the *depth* of the cortical hierarchy through which updating occurs in response to prediction error is constrained. In particular, higher gain on prediction errors drives greater sensitivity to changes in the input, promoting a model of the world that is fitted more to shorter-term fluctuations in the data. When the bottom-up input is expected to be highly precise and informative, it is less likely to be subsumed into higher-level expectations about the causes of input, compromising the inference of higher-level, context-sensitive causes of sensory input. This theory has the potential to explain a range of social, motor, cognitive and perceptual symptoms of ASD, including, for example, unusual sensitivities to sensory stimulation and reduced context sensitivity in sensory or cognitive processing (weak central coherence) (reviewed in Lawson et al., 2014; Van de Cruys et al., 2014).

### 5.2.2 Precision expectations and mentalising

Importantly, the implicit representation of others' mental states via prediction error minimization may sit at a level of the hierarchy that is particularly vulnerable to biases away from the inference of higher-level and context-sensitive causes of sensory input. As described in Section 2.3, mental states are more deeply hidden in the causal structure of the world than many other causes that we are perceptually aware of, and the relationship between mental causes and their sensory effects (including their modulation of the sensory effects of our actions) is rife with

context-sensitivity. In some respects, mentalising is the most difficult form of inference that the brain is faced with: modeling another model of the world (that itself is doing the same thing, leading to nesting and recursion). The mental states of other people modulate the sensory input in a variety of ways depending of the broader context, and hence, when modeled as expected causes, would play a higher-level contributive role in prediction rather than precisely specifying the expected sensory input in a given moment. Thus, an inflexible expectation for an overly precise mapping between causes and their sensory effects in social situations would mean that the mental states of others may not be inferred via predictive processing, nor would be modeled their contribution to the sensory consequences of potential actions. In other words, the particular complexity of the social world, necessitating greater context-sensitivity and flexibility in precision weighting for optimal learning and inference, makes the unconscious inference of mental causes of sensory data (implicit mentalising) particularly vulnerable to the hypothesised differences in precision weighting suggested to occur in ASD (Section 2.3; Hohwy & Palmer, 2014; Lawson et al., 2014; Van de Cruys et al., 2014).

An interesting development of this approach in the domain of interoception, proposed recently by Quattrocki and Friston (2014), further emphasizes the developmental trajectory of ASD together with the dynamic nature of predictive processing. Their proposal is that a primary etiological factor underlying the autistic phenotype relates to the precision weighting of interoceptive sensory signals in early child development. This is due to disruption in the function of the oxytocin system in facilitating the appropriate context-dependent gain setting for interoceptive prediction errors. In relation to social characteristics of ASD, it is suggested that an abnormality in the processing of interoceptive signals impedes the development of typical models of self and other, with a concomitant reduction in the salience of certain socially-relevant signals. This is suggested to alter the learning of adaptive social and communicative behaviour in crucial earlier stages of development. In particular, observational and imitative learning may rely on the appropriate integration of interoceptive and exteroceptive information regarding self and other, and in turn provide a basis for understanding the mental states of others.

Predictive processing accounts of ASD present a picture of neural processing in this condition as providing a model of the external world that less readily contains the mental causes of sensory input and how these causes contribute to the sensory consequences of potential actions. In contrast, ASD is often characterised as having enhanced processing or accuracy at low sensory and perceptual levels, based on a wealth of psychophysical, perceptual and cognitive data (Happé & Frith, 2006; Mottron et al., 2006; Plaisted et al., 1998). Together, this is suggestive of a developmental trajectory whereby maladaptive precision weighting at low levels precludes the formation of counterfactually-rich social-level predictive models. In the current paper we have highlighted that a lack of counterfactual predictions associated with the implicit modeling of others' mental states (i.e., active inference contingent upon others' mental states) contributes to the social challenges that individuals with ASD face, in spite of the ability of higher-functioning individuals to engage in more explicit forms of mentalising.

## 6. Conclusion

The current paper has examined mentalising as a result of the brain's inference on the (hidden) causes of its sensory input, locating mental state representations within a hierarchical cortical model of the external world. Depending on the context and timescale, mental state representations may be more or less associated with counterfactual predictions supported by active inference, thereby determining their perceptual presence. This furnishes an account of the experiential nature of implicit and explicit forms of mentalising, in terms of their associated 'mental presence', the experience of others' mental states as veridical qualities of the perceptual world. Predictive processing treatments of ASD can similarly be expanded with respect to implicit and explicit forms of mentalising and active inference. Importantly, ASD may reflect a selective deficit in the counterfactual predictions of generative models that encode deeply hidden other-mind causes of socially relevant sensory signals.

### Acknowledgements

The authors wish to thank Ryoji Sato for insightful discussions regarding perceptual presence. This work was funded by an Australian Research Council Discovery Grant (DP1311336). J.H. is supported by an Australian Research Council Future Fellowship (FT100100322). A.K.S. is grateful to the Dr. Mortimer and Theresa Sackler Foundation.

## Concluding remarks

An influential approach to the psychological study of perception is to characterise the function of the sensory system as *unconscious inference on the external causes of (ambiguous) sensory signals*, this process resulting in the representation of the world around us that we consciously experience. In the past decade, this idea has been developed in terms of *probabilistic (Bayesian) inference*, furnishing normative accounts of how the brain integrates different sources of information at an unconscious level to form an optimal model of its environment. *Predictive processing* theories cast cognitive function in terms of hierarchical probabilistic inference that is achieved as the brain makes predictions about incoming sensory data and continually works to minimise the resulting prediction error. The present thesis has explored how the neurocognitive mechanisms entailed in predictive processing shed light on the *differences* in perceptual experience and behaviour that arise between individuals and between groups, focussing on the case of ASD.

A core hypothesis that has emerged in the last several years, partly from work contained in the present thesis and partly from work by other researchers, is that ASD reflects an *overweighting of sensory signals against prior or contextual expectations* during perceptual inference. In other words, that the brain's representation of the world differs in ASD due to a systematic variation in how the sensory system integrates incoming sensory information with an existing model of the environment based on state-dependent estimates of the level of uncertainty associated with each. While much work remains in testing this hypothesis, it provides a novel perspective on a range of research findings and clinical observations, primarily those relating to sensory aspects of ASD. Existing perceptual and cognitive theories, such as the Weak Central Coherence and Enhanced Perceptual Functioning accounts (Happé & Frith, 2006; Mottron et al., 2006), have been highly successful in capturing important characteristics of ASD and motivating what is now a large body of research into non-social aspects of the condition. The attraction of a Bayesian account of ASD is partly in capturing these aspects of the condition within a framework of brain function that lends itself to

linking cognitive, computational and neural descriptions of the sensory processes involved.

Importantly, as predictive processing models have been developed in recent years to meet the challenges involved in perceptual inference, further mechanisms have been proposed that are likely to contribute to individual differences in perception and behaviour. These developments partly extend computational descriptions beyond the basic Bayesian model; specifically, the *hierarchical* nature of predictive inference, and the role of *volatility processing* in maintaining optimal inference in changing environments are critical ingredients for a rich model of perceptual inference in the brain. As argued in *Chapter 1*, these features provide needed depth to Bayesian models that better allows us to elucidate how the complex autistic phenotype could emerge from systematic differences in Bayesian mechanisms. Moreover, the extension of predictive processing to *action* is important for linking differences in perceptual inference, such as the hypothesised overweighting of sensory information in ASD, to behavioural characteristics, including those relating to motor function and the manner in which an individual samples the world around them.

Core to ASD are deficits in typical social interaction, during both early development and in adulthood. Social cognition and behaviour can similarly be explored in the context of how the brain represents the state of the world via Bayesian inference. As argued in *Chapter 6*, the brain's implicit representation of others' mental states can be cast in terms of inference on the external causes of sensory input, and in this way considered in the same framework as perceptual inference. The prominent *Theory of Mind*-deficit hypothesis of ASD can thus be considered in the light of systematic differences in how the brain infers the causes of sensory signals in ASD. Moreover, the recently proposed role for *counterfactual sensorimotor predictions* in hierarchical generative models (Friston et al., 2012; Seth, 2014, 2015a) can be extended to the social context, in which generative models include representations of others' mental states. In this way, differences in social interaction may reflect a failure to model the mental-state causes of sensory input and the way in which these modulate the sensory consequences of one's own

actions. In an active inference framework, this entails consequences for the initiation of behaviours that interact with others and the use of social actions to reduce uncertainty in perceptual models. In this way, Bayesian accounts are not just a development on theories of perception in ASD, but also suggest how differences in the processing of sensory information in the brain link to atypical behaviours more broadly.

The present thesis focussed partly on sensory processing in the context of body perception and movement (*Chapters 2 and 3*). The neuroscience of body perception has in the past decade unveiled the critical role of multisensory integration underlying the brain's representation of the body as distinct from the external environment, and in producing the phenomenological sense of ownership that we experience for our own body. The integration of seen and felt touch on the body, and the manner in which this contributes to a sense of body ownership, appears largely intact in ASD. However, autistic features in both clinical and non-clinical groups are associated with reduced sensitivity in proprioceptive and kinematic measures to illusory ownership of a fake limb (the rubber hand illusion). This coheres with the hypothesis of an increased weighting of sensory (proprioceptive) prediction errors in ASD, such that there is a reduced sensitivity to prior or contextual expectations influenced by the illusion. Moreover, adults with ASD exhibit disrupted movement initiation, which in the context of active inference is similarly suggestive of a difficulty in setting the appropriate weighting of prediction error signals against top-down predictions. These findings thus help to extend predictive processing formulations of autistic perception to motor dysfunction, which occurs commonly in ASD and, while not yet well characterised, has received increased attention in autism research in recent years (e.g., Fournier et al., 2010; Gowen & Hamilton, 2013).

The empirical results reported in the thesis provide initial evidence for certain aspects of the emerging Bayesian hypotheses for ASD, while also pointing to how these hypotheses may be refined. For instance, the intact perceptual experience of the rubber hand illusion suggest that the influence of long term expectations regarding body ownership on perception are not significantly diminished in ASD,

while altered kinematic effects of the illusion suggest that the interaction between shorter term expectations and proprioceptive feedback occurs differently in ASD compared to controls. Also examined in the present thesis was how autistic traits modulate performance on a novel statistical learning task (*Chapter 4*). This experiment provided evidence against the basic hypothesis that autistic traits are associated with a persistent overweighting of incoming information against existing expectations. This goes against the idea of a persistent difference in the weighting of sensory information (or prior expectations), which, as discussed in *Chapter 1*, might instead be developed in terms of a deficit in volatility processing. Together, these studies of multisensory processing, reach kinematics and statistical learning help to highlight that differences in the processing of sensory information may manifest in nuanced ways. Progress with Bayesian models may be made by considering factors such the operation of expectations that have developed over different timescales, and context-dependency in how sensory information is integrated with prior expectations.

Moving forward, there is an opportunity now to test Bayesian hypotheses about ASD more directly, including across psychophysical, electrophysiological and pharmacological domains. The recent emergence of techniques for modelling the role of hierarchical volatility processing in Bayesian inference provides opportunities to test finer-grained hypotheses about the specific deficit in ASD (Mathys et al., 2011; Mathys et al., 2014). Work has begun in modelling aspects of social interaction within a Bayesian framework (e.g., Sevgi, Diaconescu, Tittgemeyer, & Schilbach, 2015), and there is further scope for exploring aspects of social perception (e.g., Mareschal et al., 2013) and mentalising in a Bayesian framework for ASD; for instance, in terms of the role of state-dependent noise and volatility processing in modulating differences in ASD in these functions. An ongoing challenge will be in relating the somewhat heuristic level of description afforded by Bayesian models of brain function to the reality of the finer neural circuits that implement the processing of information in the brain. Importantly, even if we are agnostic about the veracity of Bayesian hypotheses for ASD, these theories motivate research into processes such as the role of sensory uncertainty

in modulating perceptual function that are likely to provide important insights into the nature of this condition.

There are many facets of brain function that can be related to the core activity of representing the state of the world based on ambiguous sensory data. Consequently, individual differences in the unconscious processing of sensory information in the brain potentially underlie many aspects of our conscious life and behaviour. This is a fascinating field of study because it provides a novel and computationally substantiated perspective on how the world that an individual experiences is generated in their sensory system. Moreover, this theoretical approach holds promise for shedding light on how conditions such as ASD, in which a diverse pattern of behavioural characteristics significantly impedes quality of life, emerge from systematic differences in the functioning of the nervous system. Due to the potential of Bayesian theories to span cognitive, computational and neural levels of description, this approach may be an important step towards developing medical treatments for the symptoms of ASD that cause distress (e.g., sensory hypersensitivities).

## References

- Adams, R. A., Perrinet, L. U., & Friston, K. (2012). Smooth pursuit and visual occlusion: active inference and oculomotor control in schizophrenia. *PLoS One*, 7(10), e47502. doi:10.1371/journal.pone.0047502
- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Struct Funct*, 218(3), 611-643. doi:10.1007/s00429-012-0475-5
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013). The computational anatomy of psychosis. *Front Psychiatry*, 4, 47. doi:10.3389/fpsyg.2013.00047
- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nat Neurosci*, 7(10), 1057-1058. doi:10.1038/nn1312
- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Curr Biol*, 14(3), 257-262. doi:10.1016/j.cub.2004.01.029
- American Psychiatric Association. (2000). *Diagnostic and statistical manual of mental disorders : DSM-IV-TR* (4th ed.). Washington, DC: American Psychiatric Association.
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders : DSM-5* (5th ed.). Arlington, VA: American Psychiatric Association.
- Anderson, G. M., Montazeri, F., & de Bildt, A. (2015). Network approach to autistic traits: group and subgroup analyses of ADOS item scores. *J Autism Dev Disord*, 45(10), 3115-3132. doi:10.1007/s10803-015-2537-z
- Armel, K. C., & Ramachandran, V. S. (2003). Projecting sensations to external objects: evidence from skin conductance response. *Proc Biol Sci*, 270(1523), 1499-1506. doi:10.1098/rspb.2003.2364

- Austin, E. J. (2005). Personality correlates of the broader autism phenotype as assessed by the Autism Spectrum Quotient (AQ). *Personality and Individual Differences*, 38, 451–460.
- Baranek, G. T., David, F. J., Poe, M. D., Stone, W. L., & Watson, L. R. (2006). Sensory Experiences Questionnaire: discriminating sensory features in young children with autism, developmental delays, and typical development. *J Child Psychol Psychiatry*, 47(6), 591-601. doi:10.1111/j.1469-7610.2005.01546.x
- Baron-Cohen, S. (1997). *Mindblindness: An essay on autism and theory of mind*. Cambridge, Massachusetts: MIT Press.
- Baron-Cohen, S. (2009). Autism: the empathizing-systemizing (E-S) theory. *Ann NY Acad Sci*, 1156, 68-80. doi:10.1111/j.1749-6632.2009.04467.x
- Baron-Cohen, S., Ashwin, E., Ashwin, C., Tavassoli, T., & Chakrabarti, B. (2009). Talent in autism: hyper-systemizing, hyper-attention to detail and sensory hypersensitivity. *Philos Trans R Soc Lond B Biol Sci*, 364(1522), 1377-1383. doi:10.1098/rstb.2008.0337
- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (AQ): evidence from Asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *J Autism Dev Disord*, 31(1), 5–17.
- Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nat Rev Neurosci*, 16(7), 419-429. doi:10.1038/nrn3950
- Baxter, A. J., Brugha, T. S., Erskine, H. E., Scheurer, R. W., Vos, T., & Scott, J. G. (2015). The epidemiology and global burden of autism spectrum disorders. *Psychol Med*, 45(3), 601-613. doi:10.1017/S003329171400172X
- Bayliss, A. P., & Tipper, S. P. (2005). Gaze and arrow cueing of attention reveals individual differences along the autism spectrum as a function of target context. *Br J Psychol*, 96(Pt 1), 95–114. doi:10.1348/000712604X15626

- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nat Neurosci*, 10(9), 1214-1221. doi:10.1038/nn1954
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis*, 20(3), 351-368.
- Beuker, K. T., Schjolberg, S., Kveim Lie, K., Donders, R., Lappenschaar, M., Swinkels, S. H., & Buitelaar, J. K. (2013). The structure of autism spectrum disorder symptoms in the general population at 18 months. *J Autism Dev Disord*, 43(1), 45-56. doi:10.1007/s10803-012-1546-4
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (1999). Spatio-temporal prediction modulates the perception of self-produced stimuli. *J Cogn Neurosci*, 11(5), 551-559.
- Blakemore, S. J., Tavassoli, T., Calo, S., Thomas, R. M., Catmur, C., Frith, U., & Haggard, P. (2006). Tactile sensitivity in Asperger syndrome. *Brain Cogn*, 61(1), 5-13. doi:10.1016/j.bandc.2005.12.013
- Bohl, V., & Gangopadhyay, N. (2013). Theory of mind and the unobservability of other minds. *Philosophical Explorations*, 17(2), 203-222. doi:10.1080/13869795.2013.821515
- Bölte, S., Holtmann, M., Poustka, F., Scheurich, A., & Schmidt, L. (2007). Gestalt perception and local-global processing in high-functioning autism. *J Autism Dev Disord*, 37(8), 1493-1504. doi:10.1007/s10803-006-0231-x
- Bolton, P., Macdonald, H., Pickles, A., Rios, P., Goode, S., Crowson, M., . . . Rutter, M. (1994). A case-control family history study of autism. *Journal of Child Psychology and Psychiatry*, 35(5), 877-900.
- Botvinick, M. (2004). Neuroscience. Probing the neural basis of body ownership. *Science*, 305(5685), 782-783. doi:10.1126/science.1101836

- Botvinick, M., & Cohen, J. (1998). Rubber hands 'feel' touch that eyes see. *Nature*, 391(6669), 756. doi:10.1038/35784
- Boucher, J. (2012). Putting theory of mind in its place: psychological explanations of the socio-emotional-communicative impairments in autistic spectrum disorder. *Autism*, 16(3), 226-246. doi:10.1177/1362361311430403
- Brewer, R., Happé, F., Cook, R., & Bird, G. (2015). Commentary on "Autism, oxytocin and interoception": Alexithymia, not Autism Spectrum Disorders, is the consequence of interoceptive failure. *Neurosci Biobehav Rev*, 56, 348-353. doi:10.1016/j.neubiorev.2015.07.006
- Broadbent, J., Galic, I., & Stokes, M. A. (2013). Validation of autism spectrum quotient adult version in an Australian sample. *Autism Res Treat*, 2013, 1-7. doi:10.1155/2013/984205
- Brock, J. (2012). Alternative Bayesian accounts of autistic perception: comment on Pellicano and Burr. *Trends Cogn Sci*, 16(12), 573-574; author reply 574-575. doi:10.1016/j.tics.2012.10.005
- Brock, J. (2014). Combining the old and the new: Bayesian and predictive coding accounts of autistic cognition. *The Psychologist*, 27, 750-753.
- Brock, J., Brown, C. C., Boucher, J., & Rippon, G. (2002). The temporal binding deficit hypothesis of autism. *Dev Psychopathol*, 14(2), 209-224.
- Brown, E. C., & Brune, M. (2012). The role of prediction in social neuroscience. *Front Hum Neurosci*, 6, 147. doi:10.3389/fnhum.2012.00147
- Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cogn Process*, 14(4), 411-427. doi:10.1007/s10339-013-0571-3
- Brown, H., Friston, K., & Bestmann, S. (2011). Active inference, attention, and motor preparation. *Front Psychol*, 2, 218. doi:10.3389/fpsyg.2011.00218

- Brozzoli, C., Gentile, G., & Ehrsson, H. H. (2012). That's near my hand! Parietal and premotor coding of hand-centered space contributes to localization and self-attribution of the hand. *J Neurosci*, 32(42), 14573-14582. doi:10.1523/JNEUROSCI.2660-12.2012
- Bryson, S. E., Zwaigenbaum, L., Brian, J., Roberts, W., Szatmari, P., Rombough, V., & McDermott, C. (2007). A prospective case series of high-risk infants who developed autism. *J Autism Dev Disord*, 37(1), 12-24. doi:10.1007/s10803-006-0328-2
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk : A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6(1), 3-5.
- Callenmark, B., Kjellin, L., Ronnqvist, L., & Bolte, S. (2013). Explicit versus implicit social cognition testing in autism spectrum disorder. *Autism*, 18(6), 684-693. doi:10.1177/1362361313492393
- Cascio, C. J., Foss-Feig, J. H., Burnette, C. P., Heacock, J. L., & Cosby, A. A. (2012). The rubber hand illusion in children with autism spectrum disorders: delayed influence of combined tactile and visual input on proprioception. *Autism*, 16(4), 406-419. doi:10.1177/1362361311430404
- Chakrabarti, B., Dudbridge, F., Kent, L., Wheelwright, S., Hill-Cawthorne, G., Allison, C., . . . Baron-Cohen, S. (2009). Genes related to sex steroids, neural growth, and social-emotional behavior are associated with autistic traits, empathy, and Asperger syndrome. *Autism Research*, 2(3), 157–177. doi:10.1002/aur.80
- Chen, F. S., & Yoon, J. M. (2011). Brief report: broader autism phenotype predicts spontaneous reciprocity of direct gaze. *J Autism Dev Disord*, 41(8), 1131-1134. doi:10.1007/s10803-010-1136-2
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci*, 36(3), 181-204. doi:10.1017/S0140525X12000477

- Clements, W. A., & Perner, J. (1994). Implicit understanding of belief. *Cognitive Development*, 9(4), 377-395. doi:10.1016/0885-2014(94)90012-4
- Conant, R., & Ashby, W. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89-97.
- Constantino, J. N., & Gruber, C. P. (2012). *Social Responsiveness Scale, Second Edition*. Los Angeles, CA: Western Psychological Services.
- Constantino, J. N., & Todd, R. D. (2003). Autistic traits in the general population: a twin study. *Arch Gen Psychiatry*, 60(5), 524-530. doi:10.1001/archpsyc.60.5.524
- Costantini, M., & Haggard, P. (2007). The rubber hand illusion: sensitivity and reference frame for body ownership. *Conscious Cogn*, 16(2), 229-240. doi:10.1016/j.concog.2007.01.001
- Daunizeau, J., Adam, V., & Rigoux, L. (2014). VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS Comput Biol*, 10(1), e1003441. doi:10.1371/journal.pcbi.1003441
- Davis, G., & Plaisted-Grant, K. (2015). Low endogenous neural noise in autism. *Autism*, 19(3), 351-362. doi:10.1177/1362361314552198
- Dawson, G., Toth, K., Abbott, R., Osterling, J., Munson, J., Estes, A., & Liaw, J. (2004). Early social attention impairments in autism: social orienting, joint attention, and attention to distress. *Dev Psychol*, 40(2), 271-283. doi:10.1037/0012-1649.40.2.271
- Di Martino, A., Ross, K., Uddin, L. Q., Sklar, A. B., Castellanos, F. X., & Milham, M. P. (2009). Functional brain correlates of social and nonsocial processes in autism spectrum disorders: an activation likelihood estimation meta-analysis. *Biol Psychiatry*, 65(1), 63-74. doi:10.1016/j.biopsych.2008.09.022
- Diaconescu, A. O., Mathys, C., Weber, L. A., Daunizeau, J., Kasper, L., Lomakina, E. I., . . Stephan, K. E. (2014). Inferring on the intentions of others by hierarchical

Bayesian learning. *PLoS Comput Biol*, 10(9), e1003810. doi:10.1371/journal.pcbi.1003810

Dichter, G. S., Felder, J. N., & Bodfish, J. W. (2009). Autism is characterized by dorsal anterior cingulate hyperactivation during social target detection. *Soc Cogn Affect Neurosci*, 4(3), 215-226. doi:10.1093/scan/nsp017

Dinstein, I., Heeger, D. J., Lorenzi, L., Minshew, N. J., Malach, R., & Behrmann, M. (2012). Unreliable evoked responses in autism. *Neuron*, 75(6), 981-991. doi:10.1016/j.neuron.2012.07.026

Donohue, S. E., Darling, E. F., & Mitroff, S. R. (2012). Links between multisensory processing and autism. *Experimental Brain Research*, 222(4), 377-387. doi:10.1007/s00221-012-3223-4

Dowd, A. M., McGinley, J. L., Taffe, J. R., & Rinehart, N. J. (2012). Do planning and visual integration difficulties underpin motor dysfunction in autism? A kinematic study of young children with autism. *J Autism Dev Disord*, 42(8), 1539-1548. doi:10.1007/s10803-011-1385-8

Ecker, C., Bookheimer, S. Y., & Murphy, D. G. (2015). Neuroimaging in autism spectrum disorder: brain structure and function across the lifespan. *Lancet Neurol*, 14(11), 1121-1134. doi:10.1016/S1474-4422(15)00050-2

Edwards, M. J., Adams, R. A., Brown, H., Parees, I., & Friston, K. J. (2012). A Bayesian account of 'hysteria'. *Brain*, 135(Pt 11), 3495-3512. doi:10.1093/brain/aws129

Ehrsson, H. H. (2012). The concept of body ownership and its relation to multisensory integration. In B. E. Stein (Ed.), *The new handbook of multisensory processing* (pp. xvi, 824 p., 832 p. of plates). Cambridge, Mass.: MIT Press.

Ehrsson, H. H., Holmes, N. P., & Passingham, R. E. (2005). Touching a rubber hand: feeling of body ownership is associated with activity in multisensory brain areas. *J Neurosci*, 25(45), 10564-10573. doi:10.1523/JNEUROSCI.0800-05.2005

- Ehrsson, H. H., Spence, C., & Passingham, R. E. (2004). That's my hand! Activity in premotor cortex reflects feeling of ownership of a limb. *Science*, 305(5685), 875-877. doi:10.1126/science.1097011
- Ehrsson, H. H., Wiech, K., Weiskopf, N., Dolan, R. J., & Passingham, R. E. (2007). Threatening a rubber hand that you feel is yours elicits a cortical anxiety response. *Proc Natl Acad Sci U S A*, 104(23), 9828-9833. doi:10.1073/pnas.0610011104
- Elsabbagh, M., Volein, A., Holmboe, K., Tucker, L., Csibra, G., Baron-Cohen, S., . . . Johnson, M. H. (2009). Visual orienting in the early broader autism phenotype: disengagement and facilitation. *J Child Psychol Psychiatry*, 50(5), 637-642. doi:10.1111/j.1469-7610.2008.02051.x
- Enticott, P. G., Bradshaw, J. L., Iansek, R., Tonge, B. J., & Rinehart, N. J. (2009). Electrophysiological signs of supplementary-motor-area deficits in high-functioning autism but not Asperger syndrome: an examination of internally cued movement-related potentials. *Dev Med Child Neurol*, 51(10), 787-791. doi:10.1111/j.1469-8749.2009.03270.x
- Ewbank, M. P., Rhodes, G., von dem Hagen, E. A., Powell, T. E., Bright, N., Stoyanova, R. S., . . . Calder, A. J. (2015). Repetition Suppression in Ventral Visual Cortex Is Diminished as a Function of Increasing Autistic Traits. *Cereb Cortex*, 25(10), 3381-3393. doi:10.1093/cercor/bhu149
- Ewbank, M. P., von dem Hagen, E. A., Powell, T. E., Henson, R. N., & Calder, A. J. (2015). The effect of perceptual expectation on repetition suppression to faces is not modulated by variation in autistic traits. *Cortex*. doi:10.1016/j.cortex.2015.10.011
- Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Front Hum Neurosci*, 4, 215. doi:10.3389/fnhum.2010.00215
- Field, A. (2009). *Discovering statistics using SPSS* (3rd ed.). London: Sage.

- Fiorio, M., Weise, D., Onal-Hartmann, C., Zeller, D., Tinazzi, M., & Classen, J. (2011). Impairment of the rubber hand illusion in focal hand dystonia. *Brain, 134*(Pt 5), 1428-1437. doi:10.1093/brain/awr026
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nat Rev Neurosci, 10*(1), 48-58. doi:10.1038/nrn2536
- Fournier, K. A., Hass, C. J., Naik, S. K., Lodha, N., & Cauraugh, J. H. (2010). Motor coordination in autism spectrum disorders: a synthesis and meta-analysis. *J Autism Dev Disord, 40*(10), 1227-1240. doi:10.1007/s10803-010-0981-3
- Freeth, M., Foulsham, T., & Kingstone, A. (2013). What affects social attention? Social presence, eye contact and autistic traits. *PLoS One, 8*(1), e53286. doi:10.1371/journal.pone.0053286
- Freeth, M., Sheppard, E., Ramachandran, R., & Milne, E. (2013). A cross-cultural comparison of autistic traits in the UK, India and Malaysia. *J Autism Dev Disord, 43*(11), 2569-2583. doi:10.1007/s10803-013-1808-9
- Friedman, J. (2012). How to decompose 2D trajectory data into submovements using matlab. Retrieved from <http://noisyaccumulation.blogspot.com.au/>
- Friedman, J., Brown, S., & Finkbeiner, M. (2013). Linking cognitive and reaching trajectories via intermittent movement control. *Journal of Mathematical Psychology, 57*, 140-151.
- Friston, K. (2005). A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci, 360*(1456), 815-836. doi:10.1098/rstb.2005.1622
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn Sci, 13*(7), 293-301. doi:10.1016/j.tics.2009.04.005
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat Rev Neurosci, 11*(2), 127-138. doi:10.1038/nrn2787

- Friston, K., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. *Front Psychol*, 3, 151. doi:10.3389/fpsyg.2012.00151
- Friston, K., & Frith, C. (2015). A Duet for one. *Conscious Cogn*, 36, 390-405. doi:10.1016/j.concog.2014.12.003
- Friston, K., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: the brain as a phantastic organ. *Lancet Psychiatry*, 1, 148-158.
- Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biol Cybern*, 102(3), 227-260. doi:10.1007/s00422-010-0364-z
- Friston, K. J., & Frith, C. D. (2015). Active inference, communication and hermeneutics. *Cortex*, 68, 129-143. doi:10.1016/j.cortex.2015.03.025
- Friston, K. J., Lawson, R., & Frith, C. D. (2013). On hyperpriors and hypopriors: comment on Pellicano and Burr. *Trends Cogn Sci*, 17(1), 1. doi:10.1016/j.tics.2012.11.003
- Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3), 417-458. doi:10.1007/s11229-007-9237-y
- Frith, C. (2007). *Making up the mind: how the brain creates our mental world*. Malden, MA: Blackwell Pub.
- Frith, C., & Dolan, R. J. (1997). Brain mechanisms associated with top-down processes in perception. *Philos Trans R Soc Lond B Biol Sci*, 352(1358), 1221-1230. doi:10.1098/rstb.1997.0104
- Frith, C. D., & Frith, U. (2008). Implicit and explicit processes in social cognition. *Neuron*, 60(3), 503-510. doi:10.1016/j.neuron.2008.10.032
- Frith, U. (1989). *Autism : explaining the enigma*. Oxford, UK ; Cambridge, MA, USA: Basil Blackwell.

- Frith, U. (2004). Emanuel Miller lecture: confusions and controversies about Asperger syndrome. *J Child Psychol Psychiatry*, 45(4), 672-686. doi:10.1111/j.1469-7610.2004.00262.x
- Fuentes, C. T., Mostofsky, S. H., & Bastian, A. J. (2011). No proprioceptive deficits in autism despite movement-related sensory and execution impairments. *J Autism Dev Disord*, 41(10), 1352-1361. doi:10.1007/s10803-010-1161-1
- Gallagher, S. (2004). *Philosophy, Psychiatry, & Psychology*, 11(3), 199-217.
- Gallagher, S. (2008). Direct perception in the intersubjective context. *Conscious Cogn*, 17(2), 535-543. doi:10.1016/j.concog.2008.03.003
- Geschwind, D. H., & State, M. W. (2015). Gene hunting in autism spectrum disorder: on the path to precision medicine. *Lancet Neurol*, 14(11), 1109-1120. doi:10.1016/S1474-4422(15)00044-7
- Gilbert, C. D., & Sigman, M. (2007). Brain states: top-down influences in sensory processing. *Neuron*, 54(5), 677-696. doi:10.1016/j.neuron.2007.05.019
- Glazebrook, C., Gonzalez, D., Hansen, S., & Elliott, D. (2009). The role of vision for online control of manual aiming movements in persons with autism spectrum disorders. *Autism*, 13(4), 411-433. doi:10.1177/1362361309105659
- Glazebrook, C. M., Elliott, D., & Lyons, J. (2006). A kinematic analysis of how young adults with and without autism plan and control goal-directed movements. *Motor Control*, 10(3), 244-264.
- Glazebrook, C. M., Elliott, D., & Szatmari, P. (2008). How do individuals with autism plan their movements? *J Autism Dev Disord*, 38(1), 114-126. doi:10.1007/s10803-007-0369-1
- Gómez, C., Lizier, J. T., Schaum, M., Wollstadt, P., Grützner, C., Uhlhaas, P., . . . Wibral, M. (2014). Reduced predictable information in brain signals in autism spectrum disorder. *Front Neuroinform*, 8, 9. doi:10.3389/fninf.2014.00009

- Gonzalez-Gadea, M. L., Chennu, S., Bekinschtein, T. A., Rattazzi, A., Beraudi, A., Tripicchio, P., . . . Ibanez, A. (2015). Predictive coding in autism spectrum disorder and attention deficit hyperactivity disorder. *J Neurophysiol*, 114(5), 2625-2636. doi:10.1152/jn.00543.2015
- Gowen, E., & Hamilton, A. (2013). Motor abilities in autism: a review using a computational context. *J Autism Dev Disord*, 43(2), 323-344. doi:10.1007/s10803-012-1574-0
- Green, S. A., Hernandez, L., Tottenham, N., Krasileva, K., Bookheimer, S. Y., & Dapretto, M. (2015). Neurobiology of Sensory Overresponsivity in Youth With Autism Spectrum Disorders. *JAMA Psychiatry*, 72(8), 778-786. doi:10.1001/jamapsychiatry.2015.0737
- Green, S. A., Rudie, J. D., Colich, N. L., Wood, J. J., Shirinyan, D., Hernandez, L., . . . Bookheimer, S. Y. (2013). Overreactive brain responses to sensory stimuli in youth with autism spectrum disorders. *J Am Acad Child Adolesc Psychiatry*, 52(11), 1158-1172. doi:10.1016/j.jaac.2013.08.004
- Gregory, R. L. (1980). Perceptions as hypotheses. *Philos Trans R Soc Lond B Biol Sci*, 290(1038), 181-197.
- Grinter, E. J., Maybery, M. T., Van Beek, P. L., Pellicano, E., Badcock, J. C., & Badcock, D. R. (2009). Global visual processing and self-rated autistic-like traits. *J Autism Dev Disord*, 39(9), 1278-1290. doi:10.1007/s10803-009-0740-5
- Grinter, E. J., Van Beek, P. L., Maybery, M. T., & Badcock, D. R. (2009). Brief report: visuospatial analysis and self-rated autistic-like traits. *J Autism Dev Disord*, 39(4), 670-677. doi:10.1007/s10803-008-0658-3
- Haigh, S. M., Heeger, D. J., Dinstein, I., Minshew, N., & Behrmann, M. (2015). Cortical variability in the sensory-evoked response in autism. *J Autism Dev Disord*, 45(5), 1176-1190. doi:10.1007/s10803-014-2276-6

- Happé, F., & Frith, U. (2006). The weak coherence account: detail-focused cognitive style in autism spectrum disorders. *J Autism Dev Disord*, 36(1), 5-25. doi:10.1007/s10803-005-0039-0
- Happé, F., & Ronald, A. (2008). The 'fractionable autism triad': a review of evidence from behavioural, genetic, cognitive and neural research. *Neuropsychol Rev*, 18(4), 287-304. doi:10.1007/s11065-008-9076-8
- Happé, F., Ronald, A., & Plomin, R. (2006). Time to give up on a single explanation for autism. *Nat Neurosci*, 9(10), 1218-1220. doi:10.1038/nrn1770
- Happé, F. G. (1995). The role of age and verbal ability in the theory of mind task performance of subjects with autism. *Child Dev*, 66(3), 843-855.
- Happé, F. G. (1996). Studying weak central coherence at low levels: children with autism do not succumb to visual illusions. A research note. *J Child Psychol Psychiatry*, 37(7), 873-877.
- Haswell, C. C., Izawa, J., Dowell, L. R., Mostofsky, S. H., & Shadmehr, R. (2009). Representation of internal models of action in the autistic brain. *Nat Neurosci*, 12(8), 970-972. doi:10.1038/nn.2356
- Heed, T., Grundler, M., Rinkleib, J., Rudzik, F. H., Collins, T., Cooke, E., & O'Regan, J. K. (2011). Visual information and rubber hand embodiment differentially affect reach-to-grasp actions. *Acta Psychol (Amst)*, 138(1), 263-271. doi:10.1016/j.actpsy.2011.07.003
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2), 243-259.
- Helmholtz, H. (1860). *Treatise on Physiological Optics*. New York: Dover.
- Henninger, N. A., & Taylor, J. L. (2013). Outcomes in adults with autism spectrum disorders: a historical perspective. *Autism*, 17(1), 103-116. doi:10.1177/1362361312441266

- Hill, H., & Bruce, V. (1993). Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion. *Perception*, 22(8), 887-897.
- Hill, H., & Bruce, V. (1994). A comparison between the hollow-face and 'hollow-potato' illusions. *Perception*, 23(11), 1335-1337.
- Hoekstra, R. A., Bartels, M., Cath, D. C., & Boomsma, D. I. (2008). Factor structure, reliability and criterion validity of the Autism-Spectrum Quotient (AQ): a study in Dutch population and patient groups. *J Autism Dev Disord*, 38(8), 1555-1566. doi:10.1007/s10803-008-0538-x
- Hogan, N., & Flash, T. (1987). Moving gracefully: quantitative theories of motor coordination. *Trends Neurosci*, 10(4), 170-174.
- Hogan, N., & Sternad, D. (2009). Sensitivity of smoothness measures to movement duration, amplitude, and arrests. *J Mot Behav*, 41(6), 529-534. doi:10.3200/35-09-004-RC
- Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Front Psychol*, 3, 96. doi:10.3389/fpsyg.2012.00096
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Hohwy, J. (2014). Elusive phenomenology, counterfactual awareness, and presence without mastery. *Cogn Neurosci*, 5(2), 127-128. doi:10.1080/17588928.2014.906399
- Hohwy, J., & Palmer, C. J. (2014). Social cognition as causal inference: implications for common knowledge and autism. In M. Gallotti & J. Michael (Eds.), *Social Ontology and Social Cognition* (Vol. 4): Springer.
- Hohwy, J., & Paton, B. (2010). Explaining away the body: experiences of supernaturally caused touch and touch on non-hand objects within the rubber hand illusion. *PLoS One*, 5(2), e9416. doi:10.1371/journal.pone.0009416

- Holmes, N. P., Snijders, H. J., & Spence, C. (2006). Reaching with alien limbs: visual exposure to prosthetic hands in a mirror biases proprioception without accompanying illusions of ownership. *Percept Psychophys*, 68(4), 685-701.
- Hoy, J. A., Hatton, C., & Hare, D. (2004). Weak central coherence: a cross-domain phenomenon specific to autism? *Autism*, 8(3), 267-281. doi:10.1177/1362361304045218
- Hurst, R. M., Mitchell, J. T., Kimbrel, N. A., Kwapil, T. K., & Nelson-Gray, R. O. (2007). Examination of the reliability and factor structure of the Autism Spectrum Quotient (AQ) in a non-clinical sample. *Personality and Individual Differences*, 43, 1938-1949.
- Iarocci, G., & McDonald, J. (2006). Sensory integration and the perceptual experience of persons with autism. *J Autism Dev Disord*, 36(1), 77-90. doi:10.1007/s10803-005-0044-3
- Ide, M. (2013). The effect of "anatomical plausibility" of hand angle on the rubber-hand illusion. *Perception*, 42, 103-111. doi:10.1068/p7322
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, 80(2), 519-530. doi:10.1016/j.neuron.2013.09.009
- Jakab, A., Emri, M., Spisak, T., Szeman-Nagy, A., Beres, M., Kis, S. A., . . . Berenyi, E. (2013). Autistic traits in neurotypical adults: correlates of graph theoretical functional network topology and white matter anisotropy patterns. *PLoS One*, 8(4), e60982. doi:10.1371/journal.pone.0060982
- Jarrold, C., Gilchrist, I. D., & Bender, A. (2005). Embedded figures detection in autism and typical development: preliminary evidence of a double dissociation in relationships with visual search. *Dev Sci*, 8(4), 344-351. doi:10.1111/j.1467-7687.2005.00422.x

Jeffreys, H. (1961). *The theory of probability* (3rd ed.). Oxford: Oxford University Press.

Just, M. A., Cherkassky, V. L., Keller, T. A., & Minshew, N. J. (2004). Cortical activation and synchronization during sentence comprehension in high-functioning autism: evidence of underconnectivity. *Brain*, 127(Pt 8), 1811-1821. doi:10.1093/brain/awh199

Kaldy, Z., Kraper, C., Carter, A. S., & Blaser, E. (2011). Toddlers with Autism Spectrum Disorder are more successful at visual search than typically developing toddlers. *Dev Sci*, 14(5), 980-988. doi:10.1111/j.1467-7687.2011.01053.x

Kammers, M. P., de Vignemont, F., Verhagen, L., & Dijkerman, H. C. (2009). The rubber hand illusion in action. *Neuropsychologia*, 47(1), 204-211. doi:10.1016/j.neuropsychologia.2008.07.028

Kammers, M. P., Kootker, J. A., Hogendoorn, H., & Dijkerman, H. C. (2010). How many motoric body representations can we grasp? *Exp Brain Res*, 202(1), 203-212. doi:10.1007/s00221-009-2124-7

Kammers, M. P., Longo, M. R., Tsakiris, M., Dijkerman, H. C., & Haggard, P. (2009). Specificity and coherence of body representations. *Perception*, 38(12), 1804-1820.

Kammers, M. P., Rose, K., & Haggard, P. (2011). Feeling numb: temperature, but not thermal pain, modulates feeling of body ownership. *Neuropsychologia*, 49(5), 1316-1321. doi:10.1016/j.neuropsychologia.2011.02.039

Kammers, M. P., Verhagen, L., Dijkerman, H. C., Hogendoorn, H., De Vignemont, F., & Schutter, D. J. (2009). Is this hand for real? Attenuation of the rubber hand illusion by transcranial magnetic stimulation over the inferior parietal lobule. *J Cogn Neurosci*, 21(7), 1311-1320. doi:10.1162/jocn.2009.21095

- Karaminis, T., Turi, M., Neil, L., Badcock, N. A., Burr, D., & Pellicano, E. (2015). Atypicalities in perceptual adaptation in autism do not extend to perceptual causality. *PLoS One*, 10(3), e0120439. doi:10.1371/journal.pone.0120439
- Keehn, B., Muller, R. A., & Townsend, J. (2013). Atypical attentional networks and the emergence of autism. *Neurosci Biobehav Rev*, 37(2), 164-183. doi:10.1016/j.neubiorev.2012.11.014
- Kemner, C., van Ewijk, L., van Engeland, H., & Hooge, I. (2008). Brief report: eye movements during visual search tasks indicate enhanced stimulus discriminability in subjects with PDD. *J Autism Dev Disord*, 38(3), 553-557. doi:10.1007/s10803-007-0406-0
- Kemner, C., Verbaten, M. N., Cuperus, J. M., Camfferman, G., & van Engeland, H. (1998). Abnormal saccadic eye movements in autistic children. *J Autism Dev Disord*, 28(1), 61-67.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu Rev Psychol*, 55, 271-304. doi:10.1146/annurev.psych.55.090902.142005
- Kidd, D. C., & Castano, E. (2013). Reading literary fiction improves theory of mind. *Science*, 342(6156), 377-380. doi:10.1126/science.1239918
- Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). The mirror-neuron system: a Bayesian perspective. *Neuroreport*, 18(6), 619-623. doi:10.1097/WNR.0b013e3281139ed0
- Kleinhans, N. M., Johnson, L. C., Richards, T., Mahurin, R., Greenson, J., Dawson, G., & Aylward, E. (2009). Reduced neural habituation in the amygdala and social impairments in autism spectrum disorders. *Am J Psychiatry*, 166(4), 467-475. doi:10.1176/appi.ajp.2008.07101681
- Kloosterman, P. H., Keefer, K. V., Kelley, E. A., Summerfeldt, L. J., & Parker, J. D. A. (2011). Evaluation of the factor structure of the Autism-Spectrum Quotient. *Personality and Individual Differences*, 50, 310-314.

- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci*, 27(12), 712-719. doi:10.1016/j.tins.2004.10.007
- Kok, P., & de Lange, F. P. (2014). Shape perception simultaneously up- and downregulates neural activity in the primary visual cortex. *Curr Biol*, 24(13), 1531-1535. doi:10.1016/j.cub.2014.05.042
- Kok, P., & de Lange, F. P. (2015). Predictive Coding in Sensory Cortex. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An Introduction to Model-Based Cognitive Neuroscience* (pp. 221-244). New York: Springer.
- Koster-Hale, J., & Saxe, R. (2013). Theory of mind: a neural prediction problem. *Neuron*, 79(5), 836-848. doi:10.1016/j.neuron.2013.08.020
- Krueger, J. (2012). Seeing mind in action. *Phenomenology and the Cognitive Sciences*, 11(2), 149-173. doi:10.1007/s11097-011-9226-y
- Kveraga, K., Ghuman, A. S., & Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain Cogn*, 65(2), 145-168. doi:10.1016/j.bandc.2007.06.007
- Kwakye, L. D., Foss-Feig, J. H., Cascio, C. J., Stone, W. L., & Wallace, M. T. (2011). Altered auditory and multisensory temporal processing in autism spectrum disorders. *Front Integr Neurosci*, 4, 129. doi:10.3389/fnint.2010.00129
- Lai, M. C., Lombardo, M. V., & Baron-Cohen, S. (2014). Autism. *Lancet*, 383(9920), 896-910. doi:10.1016/S0140-6736(13)61539-1
- Lai, M. C., Lombardo, M. V., Chakrabarti, B., & Baron-Cohen, S. (2013). Subgrouping the autism "spectrum": reflections on DSM-5. *PLoS Biol*, 11(4), e1001544. doi:10.1371/journal.pbio.1001544
- Lau, W. Y., Gau, S. S., Chiu, Y. N., Wu, Y. Y., Chou, W. J., Liu, S. K., & Chou, M. C. (2013). Psychometric properties of the Chinese version of the Autism Spectrum Quotient (AQ). *Res Dev Disabil*, 34(1), 294-305. doi:10.1016/j.ridd.2012.08.005

- Lau, W. Y., Kelly, A. B., & Peterson, C. C. (2013). Further Evidence on the Factorial Structure of the Autism Spectrum Quotient (AQ) for Adults With and Without a Clinical Diagnosis of Autism. *J Autism Dev Disord*, 43(12), 2807–2815. doi:10.1007/s10803-013-1827-6
- Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Front Hum Neurosci*, 8, 302. doi:10.3389/fnhum.2014.00302
- Ledesma, R. D., & Valero-Mora, P. (2007). Determining the number of factors to retain in EFA: An easy-to-use computer program for carrying out parallel analysis. *Practical Assessment, Research & Evaluation*, 12(2), 1–11.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. New York: Cambridge University Press.
- Levy, A., & Perry, A. (2011). Outcomes in adolescents and adults with autism: A review of the literature. *Research in Autism Spectrum Disorders*, 5(4), 1271–1282.
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. *Front Hum Neurosci*, 7, 547. doi:10.3389/fnhum.2013.00547
- Lloyd, D. M. (2007). Spatial limits on referred touch to an alien limb may reflect boundaries of visuo-tactile peripersonal space surrounding the hand. *Brain Cogn*, 64(1), 104-109. doi:10.1016/j.bandc.2006.09.013
- Longo, M. R., Schuur, F., Kammer, M. P., Tsakiris, M., & Haggard, P. (2008). What is embodiment? A psychometric approach. *Cognition*, 107(3), 978-998. doi:10.1016/j.cognition.2007.12.004
- Love, J., Selker, R., Marsman, M., Jamil, T., Dropmann, D., Verhagen, A. J., . . . Wagenmakers, E.-J. (2015). JASP (Version 0.7.1).
- Makin, T. R., Holmes, N. P., & Ehrsson, H. H. (2008). On the other hand: dummy hands and peripersonal space. *Behav Brain Res*, 191(1), 1-10. doi:10.1016/j.bbr.2008.02.041

- Mandy, W. P., & Skuse, D. H. (2008). Research review: What is the association between the social-communication element of autism and repetitive interests, behaviours and activities? *Journal of Child Psychology and Psychiatry*, 49(8), 795–808. doi:10.1111/j.1469-7610.2008.01911.x
- Marco, E. J., Hinkley, L. B., Hill, S. S., & Nagarajan, S. S. (2011). Sensory processing in autism: a review of neurophysiologic findings. *Pediatr Res*, 69(5 Pt 2), 48R-54R. doi:10.1203/PDR.0b013e3182130c54
- Mareschal, I., Calder, A. J., & Clifford, C. W. (2013). Humans have an expectation that gaze is directed toward them. *Curr Biol*, 23(8), 717-721. doi:10.1016/j.cub.2013.03.030
- Mari, M., Castiello, U., Marks, D., Marraffa, C., & Prior, M. (2003). The reach-to-grasp movement in children with autism spectrum disorder. *Philos Trans R Soc Lond B Biol Sci*, 358(1430), 393-403. doi:10.1098/rstb.2002.1205
- Martinez, W. L., Martinez, A. R., & Solka, J. L. (2011). *Exploratory data analysis with MATLAB* (pp. xix, 508 p., 508 p. of col).
- Masterton, B. A., & Biederman, G. B. (1983). Proprioceptive versus visual control in autistic children. *J Autism Dev Disord*, 13(2), 141-152.
- Mathys, C. D., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*, 5, 39. doi:10.3389/fnhum.2011.00039
- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Front Hum Neurosci*, 8, 825. doi:10.3389/fnhum.2014.00825
- Matson, J. L., & Williams, L. W. (2014). Depression and mood disorders among persons with Autism Spectrum Disorders. *Res Dev Disabil*, 35, 2003-2007.
- Mausfeld, R. (2013). The attribute of realness and the internal organization or perceptual reality. In L. Albertazzi (Ed.), *Handbook of Experimental*

*Phenomenology. Visual Perception of Shape, Space and Appearance.* Chichester: Wiley.

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870-881. doi:10.1016/j.neuron.2014.10.013

Michael, J., & De Bruin, L. (2015). How direct is social perception? *Conscious Cogn*, 36, 373-375. doi:10.1016/j.concog.2015.08.005

Minshew, N. J., Sung, K., Jones, B. L., & Furman, J. M. (2004). Underdevelopment of the postural control system in autism. *Neurology*, 63(11), 2056-2061.

Mitchell, P., Mottron, L., Soulieres, I., & Ropar, D. (2010). Susceptibility to the Shepard illusion in participants with autism: reduced top-down influences within perception? *Autism Res*, 3(3), 113-119. doi:10.1002/aur.130

Mitchell, P., & Ropar, D. (2004). Visuo-spatial abilities in autism: A review. *Infant and Child Development*, 13(3), 185-198.

Mongillo, E. A., Irwin, J. R., Whalen, D. H., Klaiman, C., Carter, A. S., & Schultz, R. T. (2008). Audiovisual processing in children with and without autism spectrum disorders. *J Autism Dev Disord*, 38(7), 1349-1358. doi:10.1007/s10803-007-0521-y

Moseley, G. L., Olthof, N., Venema, A., Don, S., Wijers, M., Gallace, A., & Spence, C. (2008). Psychologically induced cooling of a specific body part caused by the illusory ownership of an artificial counterpart. *Proc Natl Acad Sci U S A*, 105(35), 13169-13173. doi:10.1073/pnas.0803768105

Mottron, L., Dawson, M., Soulieres, I., Hubert, B., & Burack, J. (2006). Enhanced perceptual functioning in autism: an update, and eight principles of autistic perception. *J Autism Dev Disord*, 36(1), 27-43. doi:10.1007/s10803-005-0040-7

Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol Cybern*, 66(3), 241-251.

- Muñiz, J., García-Cueto, E., & Lozano, L. M. (2005). Item format and the psychometric properties of the Eysenck Personality Questionnaire. *Personality and Individual Differences, 38*, 61-69.
- Murray, A. L., Booth, T., McKenzie, K., Kuenssberg, R., & O'Donnell, M. (2013). Are Autistic Traits Measured Equivalently in Individuals With and Without An Autism Spectrum Disorder? An Invariance Analysis of the Autism Spectrum Quotient Short Form. *J Autism Dev Disord*. doi:10.1007/s10803-013-1851-6
- Nazarali, N., Glazebrook, C. M., & Elliott, D. (2009). Movement planning and reprogramming in individuals with autism. *J Autism Dev Disord, 39*(10), 1401-1411. doi:10.1007/s10803-009-0756-x
- Newport, R., Pearce, R., & Preston, C. (2010). Fake hands in action: embodiment and control of supernumerary limbs. *Exp Brain Res, 204*(3), 385-395. doi:10.1007/s00221-009-2104-y
- Newschaffer, C. J., Fallin, D., & Lee, N. L. (2002). Heritable and nonheritable risk factors for autism spectrum disorders. *Epidemiol Rev, 24*(2), 137-153.
- Nobile, M., Perego, P., Piccinini, L., Mani, E., Rossi, A., Bellina, M., & Molteni, M. (2011). Further evidence of complex motor dysfunction in drug naive children with autism using automatic motion analysis of gait. *Autism, 15*(3), 263-283. doi:10.1177/1362361309356929
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
- Nummenmaa, L., Engell, A. D., von dem Hagen, E., Henson, R. N., & Calder, A. J. (2012). Autism spectrum traits predict the neural response to eye gaze in typical individuals. *Neuroimage, 59*(4), 3356-3363. doi:10.1016/j.neuroimage.2011.10.075
- O'Connor, B. (2000). SPSS and SAS programs for determining the number of components using parallel analysis and Velicer's MAP test. *Behavior Research Methods, Instruments, & Computers, 32*(3), 396-402.

- O'Connor, B. (2014). SPSS, SAS, and MATLAB programs for determining the number of components and factors. <https://people.ok.ubc.ca/briocconn/nfactors/nfactors.html>. Accessed 5 August 2014.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behav Brain Sci*, 24(5), 939-973; discussion 973-1031.
- O'Riordan, M. A. (2004). Superior visual search in adults with autism. *Autism*, 8(3), 229-248. doi:10.1177/1362361304045219
- Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308(5719), 255-258. doi:10.1126/science.1107621
- Ozonoff, S., Macari, S., Young, G. S., Goldring, S., Thompson, M., & Rogers, S. J. (2008). Atypical object exploration at 12 months of age is associated with autism in a prospective sample. *Autism*, 12(5), 457-472. doi:10.1177/1362361308096402
- Palmer, C. J., Paton, B., Enticott, P. G., & Hohwy, J. (2015). 'Subtypes' in the presentation of autistic traits in the general adult population. *J Autism Dev Disord*, 45(5), 1291-1301. doi:10.1007/s10803-014-2289-1
- Palmer, C. J., Paton, B., Hohwy, J., & Enticott, P. G. (2013). Movement under uncertainty: the effects of the rubber-hand illusion vary along the nonclinical autism spectrum. *Neuropsychologia*, 51(10), 1942-1951. doi:10.1016/j.neuropsychologia.2013.06.020
- Palmer, C. J., Paton, B., Kirkovski, M., Enticott, P. G., & Hohwy, J. (2015). Context sensitivity in action decreases along the autism spectrum: a predictive processing perspective. *Proc Biol Sci*, 282(1802). doi:10.1098/rspb.2014.1557
- Palmer, C. J., Seth, A. K., & Hohwy, J. (2015). The felt presence of other minds: Predictive processing, counterfactual predictions, and mentalising in autism. *Conscious Cogn*, 36, 376-389. doi:10.1016/j.concog.2015.04.007

- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5), 411–419.
- Paton, B., Hohwy, J., & Enticott, P. G. (2012). The rubber hand illusion reveals proprioceptive and sensorimotor differences in autism spectrum disorders. *J Autism Dev Disord*, 42(9), 1870-1883. doi:10.1007/s10803-011-1430-7
- Pavani, F., Spence, C., & Driver, J. (2000). Visual capture of touch: out-of-the-body experiences with rubber gloves. *Psychol Sci*, 11(5), 353-359.
- Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput Biol*, 7(1), e1001048. doi:10.1371/journal.pcbi.1001048
- Pellicano, E., & Burr, D. (2012). When the world becomes 'too real': a Bayesian explanation of autistic perception. *Trends Cogn Sci*, 16(10), 504-510. doi:10.1016/j.tics.2012.08.009
- Pellicano, E., Jeffery, L., Burr, D., & Rhodes, G. (2007). Abnormal adaptive face-coding mechanisms in children with autism spectrum disorder. *Curr Biol*, 17(17), 1508-1512. doi:10.1016/j.cub.2007.07.065
- Pelphrey, K. A., Shultz, S., Hudac, C. M., & Vander Wyk, B. C. (2011). Research review: Constraining heterogeneity: the social brain and its development in autism spectrum disorder. *J Child Psychol Psychiatry*, 52(6), 631-644. doi:10.1111/j.1469-7610.2010.02349.x
- Petkova, V. I., & Ehrsson, H. H. (2008). If I were you: perceptual illusion of body swapping. *PLoS One*, 3(12), e3832. doi:10.1371/journal.pone.0003832
- Petkova, V. I., & Ehrsson, H. H. (2009). When right feels left: referral of touch and ownership between the hands. *PLoS One*, 4(9), e6933. doi:10.1371/journal.pone.0006933
- Piven, J., Palmer, P., Jacobi, D., Childress, D., & Arndt, S. (1997). Broader autism phenotype: evidence from a family history study of multiple-incidence autism families. *American Journal of Psychiatry*, 154(2), 185–190.

- Plaisted, K., O'Riordan, M., & Baron-Cohen, S. (1998). Enhanced discrimination of novel, highly similar stimuli by adults with autism during a perceptual learning task. *J Child Psychol Psychiatry*, 39(5), 765-775.
- Posserud, M. B., Lundervold, A. J., & Gillberg, C. (2006). Autistic features in a total population of 7-9-year-old children assessed by the ASSQ (Autism Spectrum Screening Questionnaire). *Journal of Child Psychology and Psychiatry*, 47(2), 167-175. doi:10.1111/j.1469-7610.2005.01462.x
- Puts, N. A., Wodka, E. L., Tommerdahl, M., Mostofsky, S. H., & Edden, R. A. (2014). Impaired tactile processing in children with autism spectrum disorder. *J Neurophysiol*, 111(9), 1803-1811. doi:10.1152/jn.00890.2013
- Puzzo, I., Cooper, N. R., Vetter, P., & Russo, R. (2010). EEG activation differences in the pre-motor cortex and supplementary motor area between normal individuals with high and low traits of autism. *Brain Res*, 1342, 104-110. doi:10.1016/j.brainres.2010.04.060
- Qian, N., & Lipkin, R. M. (2011). A learning-style theory for understanding autistic behaviors. *Front Hum Neurosci*, 5, 77. doi:10.3389/fnhum.2011.00077
- Quattrocki, E., & Friston, K. (2014). Autism, oxytocin and interoception. *Neurosci Biobehav Rev*, 47, 410-430. doi:10.1016/j.neubiorev.2014.09.012
- Quoidbach, J., Gilbert, D. T., & Wilson, T. D. (2013). The end of history illusion. *Science*, 339(6115), 96-98. doi:10.1126/science.1229294
- Rinehart, N. J., Bellgrove, M. A., Tonge, B. J., Brereton, A. V., Howells-Rankin, D., & Bradshaw, J. L. (2006). An examination of movement kinematics in young people with high-functioning autism and Asperger's disorder: further evidence for a motor planning deficit. *J Autism Dev Disord*, 36(6), 757-767. doi:10.1007/s10803-006-0118-x
- Rinehart, N. J., Bradshaw, J. L., Brereton, A. V., & Tonge, B. J. (2001). Movement preparation in high-functioning autism and Asperger disorder: a serial choice

- reaction time task involving motor reprogramming. *J Autism Dev Disord*, 31(1), 79-88.
- Rinehart, N. J., Tonge, B. J., Bradshaw, J. L., Iansek, R., Enticott, P. G., & Johnson, K. A. (2006). Movement-related potentials in high-functioning autism and Asperger's disorder. *Dev Med Child Neurol*, 48(4), 272-277. doi:10.1017/S0012162206000594
- Ring, H., Woodbury-Smith, M., Watson, P., Wheelwright, S., & Baron-Cohen, S. (2008). Clinical heterogeneity among people with high functioning autism spectrum conditions: evidence favouring a continuous severity gradient. *Behavioral and Brain Functions*, 4, 11. doi:10.1186/1744-9081-4-11
- Ring, H. A., Baron-Cohen, S., Wheelwright, S., Williams, S. C., Brammer, M., Andrew, C., & Bullmore, E. T. (1999). Cerebral correlates of preserved cognitive skills in autism: a functional MRI study of embedded figures task performance. *Brain*, 122 (Pt 7), 1305-1315.
- Ritvo, R. A., Ritvo, E. R., Guthrie, D., Ritvo, M. J., Hufnagel, D. H., McMahon, W., . . . Eloff, J. (2011). The Ritvo Autism Asperger Diagnostic Scale-Revised (RAADS-R): a scale to assist the diagnosis of Autism Spectrum Disorder in adults: an international validation study. *J Autism Dev Disord*, 41(8), 1076-1089. doi:10.1007/s10803-010-1133-5
- Robertson, A. E., & Simmons, D. R. (2013). The relationship between sensory sensitivity and autistic traits in the general population. *J Autism Dev Disord*, 43(4), 775-784. doi:10.1007/s10803-012-1608-7
- Robic, S., Sonie, S., Fonlupt, P., Henaff, M. A., Touil, N., Coricelli, G., . . . Schmitz, C. (2015). Decision-making in a changing world: a study in autism spectrum disorders. *J Autism Dev Disord*, 45(6), 1603-1613. doi:10.1007/s10803-014-2311-7
- Robinson, E. B., Koenen, K. C., McCormick, M. C., Munir, K., Hallett, V., Happé, F., . . . Ronald, A. (2011). Evidence that autistic traits show the same etiology in the

- general population and at the quantitative extremes (5%, 2.5%, and 1%). *Arch Gen Psychiatry*, 68(11), 1113–1121. doi:10.1001/archgenpsychiatry.2011.119
- Robinson, E. B., Koenen, K. C., McCormick, M. C., Munir, K., Hallett, V., Happé, F., . . . Ronald, A. (2012). A multivariate twin study of autistic traits in 12-year-olds: testing the fractionable autism triad hypothesis. *Behav Genet*, 42(2), 245–255. doi:10.1007/s10519-011-9500-3
- Rogers, S. J., & Ozonoff, S. (2005). Annotation: what do we know about sensory dysfunction in autism? A critical review of the empirical evidence. *J Child Psychol Psychiatry*, 46(12), 1255–1268. doi:10.1111/j.1469-7610.2005.01431.x
- Rohde, M., Di Luca, M., & Ernst, M. O. (2011). The Rubber Hand Illusion: feeling of ownership and proprioceptive drift do not go hand in hand. *PLoS One*, 6(6), e21659. doi:10.1371/journal.pone.0021659
- Romero, D. H., Van Gemmert, A. W., Adler, C. H., Bekkering, H., & Stelmach, G. E. (2003). Altered aiming movements in Parkinson's disease patients and elderly adults as a function of delays in movement onset. *Exp Brain Res*, 151(2), 249–261. doi:10.1007/s00221-003-1452-2
- Ronald, A., Happé, F., Bolton, P., Butcher, L. M., Price, T. S., Wheelwright, S., . . . Plomin, R. (2006). Genetic heterogeneity between the three components of the autism spectrum: a twin study. *J Am Acad Child Adolesc Psychiatry*, 45(6), 691–699. doi:10.1097/01.chi.0000215325.13058.9d
- Ronald, A., Happé, F., & Plomin, R. (2005). The genetic relationship between individual differences in social and nonsocial behaviours characteristic of autism. *Dev Sci*, 8(5), 444–458. doi:10.1111/j.1467-7687.2005.00433.x
- Ronald, A., Happé, F., & Plomin, R. (2006). Genetic research into autism. *Science*, 311(5763), 952; author reply 952. doi:10.1126/science.311.5763.952a
- Ronald, A., Larsson, H., Anckarsater, H., & Lichtenstein, P. (2011). A twin study of autism symptoms in Sweden. *Mol Psychiatry*, 16(10), 1039–1047. doi:10.1038/mp.2010.82

- Ropar, D., & Mitchell, P. (1999). Are individuals with autism and Asperger's syndrome susceptible to visual illusions? *J Child Psychol Psychiatry*, 40(8), 1283-1293.
- Ropar, D., & Mitchell, P. (2001). Susceptibility to illusions and performance on visuospatial tasks in individuals with autism. *J Child Psychol Psychiatry*, 42(4), 539-549.
- Sasson, N. J., Elison, J. T., Turner-Brown, L. M., Dichter, G. S., & Bodfish, J. W. (2011). Brief report: Circumscribed attention in young children with autism. *J Autism Dev Disord*, 41(2), 242-247. doi:10.1007/s10803-010-1038-3
- Sasson, N. J., Turner-Brown, L. M., Holtzclaw, T. N., Lam, K. S., & Bodfish, J. W. (2008). Children with autism demonstrate circumscribed attention during passive viewing of complex social and nonsocial picture arrays. *Autism Res*, 1(1), 31-42. doi:10.1002/aur.4
- Schilbach, L. (2014). On the relationship of online and offline social cognition. *Front Hum Neurosci*, 8, 278. doi:10.3389/fnhum.2014.00278
- Schilbach, L. (2015). Eye to eye, face to face and brain to brain: novel approaches to study the behavioral dynamics and neural mechanisms of social interactions. *Current Opinion in Behavioral Sciences*, 3, 130-135.
- Schilbach, L. (2016). Towards a second-person neuropsychiatry. *Philos Trans R Soc Lond B Biol Sci*, 371(1686). doi:10.1098/rstb.2015.0081
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., & Vogeley, K. (2013). Toward a second-person neuroscience. *Behav Brain Sci*, 36(4), 393-414. doi:10.1017/S0140525X12000660
- Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., . . . Vogeley, K. (2010). Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *J Cogn Neurosci*, 22(12), 2702-2715. doi:10.1162/jocn.2009.21401

- Scholl, B. J., & Gao, T. (2013). Perceiving animacy and intentionality: Visual processing or higher-level judgment? In M. D. Rutherford & V. A. Kuhlmeier (Eds.), *Social perception: Detection and interpretation of animacy, agency, and intention* (pp. 197-230). Cambridge, MA: MIT Press.
- Senju, A., & Johnson, M. H. (2009). Atypical eye contact in autism: models, mechanisms and development. *Neurosci Biobehav Rev*, 33(8), 1204-1214. doi:10.1016/j.neubiorev.2009.06.001
- Senju, A., Southgate, V., White, S., & Frith, U. (2009). Mindblind eyes: an absence of spontaneous theory of mind in Asperger syndrome. *Science*, 325(5942), 883-885. doi:10.1126/science.1176170
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends Cogn Sci*, 17(11), 565-573. doi:10.1016/j.tics.2013.09.007
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cogn Neurosci*, 5(2), 97-118. doi:10.1080/17588928.2013.877880
- Seth, A. K. (2015a). The cybernetic Bayesian brain: from interoceptive inference to sensorimotor contingencies. In T. Metzinger & J. M. Windt (Eds.), *Open MIND* (Vol. 35(T)). Frankfurt am Main: MIND Group.
- Seth, A. K. (2015b). Presence, objecthood, and the phenomenology of predictive perception. *Cogn Neurosci*, 6(2-3), 111-117. doi:10.1080/17588928.2015.1026888
- Sevgi, M., Diaconescu, A. O., Tittgemeyer, M., & Schilbach, L. (2015). Social Bayes: Using Bayesian Modeling to Study Autistic Trait-Related Differences in Social Cognition. *Biol Psychiatry*. doi:10.1016/j.biopsych.2015.11.025
- Shah, A., & Frith, U. (1983). An islet of ability in autistic children: a research note. *J Child Psychol Psychiatry*, 24(4), 613-620.
- Shimada, S., Fukuda, K., & Hiraki, K. (2009). Rubber hand illusion under delayed visual feedback. *PLoS One*, 4(7), e6185. doi:10.1371/journal.pone.0006185

- Shipp, S., Adams, R. A., & Friston, K. J. (2013). Reflections on agranular architecture: predictive coding in the motor cortex. *Trends Neurosci*, 36(12), 706-716. doi:10.1016/j.tins.2013.09.004
- Shuster, J., Perry, A., Bebko, J., & Toplak, M. E. (2013). Review of Factor Analytic Studies Examining Symptoms of Autism Spectrum Disorders. *J Autism Dev Disord*. doi:10.1007/s10803-013-1854-3
- Silberman, S. (2015). *NeuroTribes: The Legacy of Autism and the Future of Neurodiversity*. St Leonards, Australia: Allen & Unwin.
- Simmons, D. R., Robertson, A. E., McKay, L. S., Toal, E., McAleer, P., & Pollick, F. E. (2009). Vision in autism spectrum disorders. *Vision Res*, 49(22), 2705-2739. doi:10.1016/j.visres.2009.08.005
- Skewes, J. C., Jegindo, E. M., & Gebauer, L. (2014). Perceptual inference and autistic traits. *Autism*. doi:10.1177/1362361313519872
- Smith, J. (2010). Seeing Other People. *Philosophy and Phenomenological Research*, 81(3), 731-748. doi:10.1111/j.1933-1592.2010.00392.x
- Stevens, J. (2009). *Applied multivariate statistics for the social sciences* (5th ed.). New York: Routledge.
- Stewart, M. E., & Austin, E. J. (2009). The structure of the Autism-Spectrum Quotient (AQ): Evidence from a student sample in Scotland. *Personality and Individual Differences*, 47, 224–228.
- Stewart, M. E., Griffiths, T. D., & Grube, M. (2015). Autistic Traits and Enhanced Perceptual Representation of Pitch and Time. *J Autism Dev Disord*. doi:10.1007/s10803-015-2517-3
- Stewart, M. E., Watson, J., Allcock, A. J., & Yaqoob, T. (2009). Autistic traits predict performance on the block design. *Autism*, 13(2), 133-142. doi:10.1177/1362361308098515

- Stoit, A. M., van Schie, H. T., Slaats-Willemse, D. I., & Buitelaar, J. K. (2013). Grasping motor impairments in autism: not action planning but movement execution is deficient. *J Autism Dev Disord*, 43(12), 2793-2806. doi:10.1007/s10803-013-1825-8
- Summerfield, C., Tritschuh, E. H., Monti, J. M., Mesulam, M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nat Neurosci*, 11(9), 1004-1006. doi:10.1038/nn.2163
- Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychol Sci*, 18(7), 580-586. doi:10.1111/j.1467-9280.2007.01943.x
- Sutherland, A., & Crewther, D. P. (2010). Magnocellular visual evoked potential delay with high autism spectrum quotient yields a neural mechanism for altered perception. *Brain*, 133(Pt 7), 2089–2097. doi:10.1093/brain/awq122
- Tager-Flusberg, H. (2001). A Reexamination of the Theory of Mind Hypothesis of Autism. In J. A. Burack, T. Charman, N. Yirmiya, & P. R. Zelazo (Eds.), *The Development of Autism: Perspectives From Theory and Research*. New Jersey: Lawrence Erlbaum Associates.
- Teulings, H. L., Contreras-Vidal, J. L., Stelmach, G. E., & Adler, C. H. (1997). Parkinsonism reduces coordination of fingers, wrist, and arm in fine motor control. *Exp Neurol*, 146(1), 159-170. doi:10.1006/exnr.1997.6507
- Thakkar, K. N., Nichols, H. S., McIntosh, L. G., & Park, S. (2011). Disturbances in body ownership in schizophrenia: evidence from the rubber hand illusion and case study of a spontaneous out-of-body experience. *PLoS One*, 6(10), e27089. doi:10.1371/journal.pone.0027089
- Thakkar, K. N., Polli, F. E., Joseph, R. M., Tuch, D. S., Hadjikhani, N., Barton, J. J., & Manoach, D. S. (2008). Response monitoring, repetitive behaviour and anterior cingulate abnormalities in autism spectrum disorders (ASD). *Brain*, 131(Pt 9), 2464-2478. doi:10.1093/brain/awn099

- Thinès, G., Costall, A., & Butterworth, G. (Eds.). (1991). *Michotte's experimental phenomenology of perception*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Tibshirani, R., Walther, G., & Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B*, 63(2), 411–423.
- Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *J Neurosci*, 31(25), 9118-9123. doi:10.1523/JNEUROSCI.1425-11.2011
- Tsakiris, M., Carpenter, L., James, D., & Fotopoulou, A. (2010). Hands only illusion: multisensory integration elicits sense of ownership for body parts but not for non-corporeal objects. *Exp Brain Res*, 204(3), 343-352. doi:10.1007/s00221-009-2039-3
- Tsakiris, M., & Haggard, P. (2005). The rubber hand illusion revisited: visuotactile integration and self-attribution. *J Exp Psychol Hum Percept Perform*, 31(1), 80-91. doi:10.1037/0096-1523.31.1.80
- Turi, M., Burr, D. C., Igliozi, R., Aagten-Murphy, D., Muratori, F., & Pellicano, E. (2015). Children with autism spectrum disorder show reduced adaptation to number. *Proc Natl Acad Sci U S A*, 112(25), 7868-7872. doi:10.1073/pnas.1504099112
- Turner, K. C., Frost, L., Linsenbardt, D., McIlroy, J. R., & Muller, R. A. (2006). Atypically diffuse functional connectivity between caudate nuclei and cerebral cortex in autism. *Behav Brain Funct*, 2, 34. doi:10.1186/1744-9081-2-34
- van Beers, R. J., Sittig, A. C., & Gon, J. J. (1998). The precision of priorceptive position sense. *Exp Brain Res*, 122, 367–377.
- van Beers, R. J., Sittig, A. C., & Gon, J. J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *J Neurophysiol*, 81(3), 1355-1364.

- van Beers, R. J., Wolpert, D. M., & Haggard, P. (2002). When feeling is more important than seeing in sensorimotor adaptation. *Curr Biol*, 12(10), 834-837.
- van Boxtel, J. J., & Lu, H. (2013a). Impaired global, and compensatory local, biological motion processing in people with high levels of autistic traits. *Front Psychol*, 4, 209. doi:10.3389/fpsyg.2013.00209
- van Boxtel, J. J., & Lu, H. (2013b). A predictive coding perspective on autism spectrum disorders. *Front Psychol*, 4, 19. doi:10.3389/fpsyg.2013.00019
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*.
- Van Doorn, G., Hohwy, J., & Symmons, M. (2014). Can you tickle yourself if you swap bodies with someone else? *Conscious Cogn*, 23, 1-11. doi:10.1016/j.concog.2013.10.009
- Vilares, I., & Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann N Y Acad Sci*, 1224, 22-39. doi:10.1111/j.1749-6632.2011.05965.x
- Walter, E., Dassonville, P., & Bochsler, T. M. (2009). A specific autistic trait that modulates visuospatial illusion susceptibility. *J Autism Dev Disord*, 39(2), 339-349. doi:10.1007/s10803-008-0630-2
- Wass, S. V., Jones, E. J., Gliga, T., Smith, T. J., Charman, T., Johnson, M. H., & team, B. (2015). Shorter spontaneous fixation durations in infants with later emerging autism. *Sci Rep*, 5, 8284. doi:10.1038/srep08284
- Werner, E., Dawson, G., Munson, J., & Osterling, J. (2005). Variation in early developmental course in autism and its relation with behavioral outcome at 3-4 years of age. *J Autism Dev Disord*, 35(3), 337-350.
- White, R. C., & Aimola Davies, A. M. (2011). Touching my left elbow: the anatomical structure of the body affects the illusion of self-touch. *Perception*, 40(1), 95-98.

- World Health Organization. (2004). *ICD-10 : international statistical classification of diseases and related health problems* (10th revision, 2nd ed.). Geneva: World Health Organization.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4), 681-692. doi:10.1016/j.neuron.2005.04.026
- Zablotsky, B., Black, L., Maenner, M. J., Schieve, L. A., & Blumberg, S. J. (2015). Estimated Prevalence of Autism and Other Developmental Disabilities Following Questionnaire Changes in the 2014 National Health Interview Survey. *National Health Statistics Report*, 2015, 87, 1-20.
- Zahavi, D. (2011). Empathy and Direct Social Perception: A Phenomenological Proposal. *Review of Philosophy and Psychology*, 2(3), 541-558. doi:10.1007/s13164-011-0070-3
- Zaidel, A., Goin-Kochel, R. P., & Angelaki, D. E. (2015). Self-motion perception in autism is compromised by visual noise but integrated optimally across multiple senses. *Proc Natl Acad Sci U S A*, 112(20), 6461-6466. doi:10.1073/pnas.1506582112
- Zaki, J. (2013). Cue Integration: A Common Framework for Social Cognition and Physical Perception. *Perspectives on Psychological Science*, 8(3), 296-312. doi:10.1177/1745691613475454
- Zopf, R., Savage, G., & Williams, M. A. (2010). Crossmodal congruency measures of lateral distance effects on the rubber hand illusion. *Neuropsychologia*, 48(3), 713-725. doi:10.1016/j.neuropsychologia.2009.10.028
- Zopf, R., Truong, S., Finkbeiner, M., Friedman, J., & Williams, M. A. (2011). Viewing and feeling touch modulates hand position for reaching. *Neuropsychologia*, 49(5), 1287-1293. doi:10.1016/j.neuropsychologia.2011.02.012
- Zwaigenbaum, L., Bryson, S., Lord, C., Rogers, S., Carter, A., Carver, L., . . . Yirmiya, N. (2009). Clinical assessment and management of toddlers with suspected

autism spectrum disorder: insights from studies of high-risk infants. *Pediatrics*, 123(5), 1383-1391. doi:10.1542/peds.2008-1606

Zwaigenbaum, L., Bryson, S., Rogers, T., Roberts, W., Brian, J., & Szatmari, P. (2005). Behavioral manifestations of autism in the first year of life. *Int J Dev Neurosci*, 23(2-3), 143-152. doi:10.1016/j.ijdevneu.2004.05.001

# Appendices

1. Published versions of papers included in the main text of the thesis
  - a. [Chapter 2] Palmer, CJ, Paton, B, Hohwy, J, Enticott, P. (2013). Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum. *Neuropsychologia*, 51(10), 1942–1951.
  - b. [Chapter 3] Palmer, CJ, Paton, B, Kirkovski, M, Enticott, PG, Hohwy, J. (2015). Context sensitivity in action decreases along the autism spectrum: a predictive processing perspective. *Proceedings of the Royal Society B: Biological Sciences*, 282(1802).
  - c. [Chapter 5] Palmer, CJ, Paton, B, Enticott, P, Hohwy, J. (2015). ‘Subtypes’ in the presentation of autistic traits in the general adult population. *Journal of Autism and Developmental Disorders*, 45(5), 1291–1301.
  - d. [Chapter 6] Palmer, CJ, Seth, A, Hohwy, J. (2015). The felt presence of other minds: predictive processing, counterfactual predictions, and mentalising in autism. *Consciousness & Cognition*, 36, 376–389.
2. Further appendices
  - a. Hohwy, J, Palmer, C. (2014). Social cognition as causal inference: implications for common knowledge and autism. In Mattia Gallotti and John Michael (Eds.), *Social Ontology and Social Cognition*, Springer Series “Studies in the Philosophy of Sociality”, Vol. 4, 2014.