MAKING YOUR ACQUAINTANCE:

PERSPECTIVAL EXPERIENCE AS SENSORIMOTOR ACCOMPLISHMENT

Andrew George Atkinson Britten-Neish

(Student number: 12009287)

UCL

PhD

Declaration

I, Andrew George Atkinson Britten-Neish, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Andrew George Atkinson Britten-Neish

Abstract

My thesis argues for a relational account of the content and structure of perceptual experience, the sensorimotor acquaintance view. This claims that a core class of perceptual episodes, perspectival experiences, are grounded in the exercise of sensorimotor capacities on environmental objects and the obtaining of representational states in the brain. Relational models of perception are at the centre of two approaches in contemporary philosophy of mind and cognitive science that are typically opposed to representational explanation. Naïve realists argue that perception is a 'personal-level' relation of acquaintance with environmental objects; enactivists claim that the 'subpersonal basis of perception involve similar concrete relations between perceivers and their environment. The sensorimotor view brings together these two strands of thought, drawing on the work of Susan Hurley (a prominent influence on the work of many contemporary enactivists), whose account of how personal and subpersonal levels relate is a crucial feature of the view. Setting Hurley's work in a context that includes many of the philosophical influences on naïve realism, I argue that the central claims naïve realists make about perceptual experience are best articulated through this account.

Impact Statement

The thesis intervenes in several ongoing debates within the philosophy of cognitive science. It articulates a novel account of the distinction between personal and subpersonal levels of explanation, around which interdisciplinary discussion about the consequences that mechanistic explanation in the brain sciences has for the wider scientific and humanistic study of the mind. In particular, it deals with the relation between models of perceptual experience articulated by philosophers and neurocomputational mechanisms for object perception within two contemporary paradigms: 'relationist' models of perception and the 'object file' construct in perceptual psychology and neuroscience. It supports its conclusions with an interpretation of the work of two influential interdisciplinary philosophers of cognitive science, Dan Dennett and Susan Hurley, and contributes to a growing reception of the latter's work in contemporary philosophy.

Table of Contents

ACKNOWLEDGEMENTS 7

INTRODUCTION 8

PART I: PERCEPTUAL ACQUAINTANCE AND LEVELS OF EXPLANATION

1. Naïve Realism, Neurocomputation and Perceptual Acquaintance

1.0 Introduction 20

1.1 Acquaintance, transparency and realism about the objects of perceptual Experience 21

1.2 Naïve realism and other contemporary views 34

1.3 Neurocomputational naïve realism 42

1.4 Conclusion 55

2. Hurley, Dennett and the Philosophers of the Personal Level: Psychological Interpretation and Subpersonal Explanation

2.0 Introduction 57

2.1 Psychological Interpretation and the Personal/Subpersonal Distinction 60

2.2 Interpretation and the Philosophers of the Personal Level 73

2.3 Subpersonal Explanation: What to Do when Your Spade is Turned 79

2.4 Neurocomputational Mechanisms and Isomorphic Explanation 85

2.5 Isomorphic and Non-Isomorphic Styles of Subpersonal Explanation 91

2.6 Conclusion 103

3. Agents, their Actions and the Perspectival Unity of Experience

3.0 Introduction 104

3.1 Isomorphic Content and Perspectival Unity 106

3.2 The Standard Story of Action and Isomorphic Explanation 117

3.3 Perspectival Unity and Basic Actions 124

3.4 Basic Particulars in an Action Space 131

3.5 Conclusion 141

PART II: THE SENSORIMOTOR ACQUAINTANCE VIEW

4. Objects of Perception and Environmental Objects

4.0 Introduction 143

4.1 Appearances and Perceptual Objects 145

4.2 Perceptual Objects 153

4.3 Object Files 159

4.4 Conclusion 166

5. Accomplishment and Acquaintance

5.0 Introduction 168

5.1 Acquaintance as Epistemic Skill: Dickie's Semantic Acquaintance Theory 169

5.2 Object Files and the Subpersonal Conditions of Acquaintance 184

5.3 Getting Acquainted: Agential Skill and Perceptual Objects 194

5.4 Conclusion 204

BIBLIOGRAPHY 205

Acknowledgments

First of all, I would like to gratefully acknowledge help given over the course of my PhD by three supervisors, Lucy O'Brien, Doug Lavin and, especially, Mark Kalderon. Mark likes to say that existence is the first virtue of a thesis. It's a virtue that were it not for his patient guidance during the crucial last stages of its production, this one might well have lacked. Thanks for this and more are also due to the UCL graduate community. For encouragement, discussion and comments on one or other part of the project, these particularly go to Julian Bacharach, Léa Bourguignon, Kirstine La Cour, Harry Long, Michael Marcunas, Alex Murphy, Rory Philips, Filippa Ronquist, Pablo Fernandez Velasco and Nikhil Venkatesh. Further afield, heartfelt thanks are owed to Zoe Drayson and Rob Rupert for comments on what became chapters 1 and 2 and to Dave Ward and Julian Kiverstein for inspiration and advice on material related to Susan Hurley. I would like also to thank my examiners Adrian Alsmith and Alex Grzankowski for the warm intellectual atmosphere they created during the viva and the generous spirit in which they engaged with what I wrote (or intended to write). My parents, Sara and Andrew, have encouraged me from my first taking an interest in philosophy – something I'm very grateful for and whose results I hope they're pleased with. Finally, to my husband Robbie: I can only again say here how outrageously lucky I feel – for your support in the last five years while I was doing this and for everything else now that it's done. Thank you.

Introduction

a) A hackneyed topic

In June of 1951, C. D. Broad opened the lecture later published as *Some Elementary Reflections on Sense Perception* (1952) with an apology for the "hackneyed topic" he had chosen for discussion. He continued, noting that 'perception', or "sense-perception", is a technical term – though, he hastened to add, "none the worse for that" (p. 3) – used by philosophers and psychologists when discussing "the experiences which we describe in daily life as "seeing" "hearing" "tasting" "touching" "smelling" and perhaps some others" (ibid.). By describing 'perception' as a technical term, Broad had in mind more than the fact that the specific terms he mentioned are more common or familiar than the generic 'perceiving'. Technical terms, or familiar terms used in their 'technical senses', get their meaning from their place in a wider technical vocabulary. Such vocabularies contain sets of terms whose meaning is restricted to their application to a field of study, practice or discourse. Broad assumed in his audience a deep, even over-, familiarity with the background discourse against which 'sense-perception' was already (in 1951) both a technical term and hackneyed topic. This background is not the background that informed the great majority of the historical discussions now retrospectively classed as the 'philosophy of perception' from the pre-Socratics through to the neo-Hegelians. Nor, I'll suggest below, is it the background that contemporary discussion of perception in the philosophy of mind and cognitive science *fully* share in. Broad dates the relevant period and discourse by introducing another core element of its terminology, "the notion of "sense-data" […] which played so great a part in the philosophy of sense-perception in the first thirty years of this century and has been so heavily belaboured since then" (p. 12).

'Sense-data' is a term that wears its technical nature on its sleeve and, as Broad notes, had played a central organising role in discussion of perception in analytic philosophy since the earliest years of the emerging movement. By the time he was giving his lecture, however, this central role was coming under pressure – from the 'ordinary language philosophers', Wittgenstein and Austin – and would soon be further undermined by different kinds of 'intentionalist' theories drawing on this tradition, (e.g., Anscombe 1965) or transcribing Kantian ideas about the subject and object of apperception into a style of thinking shaped by it (Sellars 1956; Strawson 1966). Since then, it has fallen (mostly) out of favour – but not out of use. Philosophers continue to talk about sense data, even though sense data analyses of perceptual phenomena often play something more like the role of 'Cartesian' solutions to the mind-body

problem. Their rejection rather than their widespread acceptance is more likely to make their evocation hackneyed.

Especially since the 'cognitive revolution' in philosophy and the mind sciences, the 'representational theory of mind' has been used to support theories of perception that are also frequently opposed to 'sense data accounts' (see, e.g. Pautz 2021a). But it is far from clear that the sense data theories with which 'representationalism' is contrasted are commitment to the same set of presuppositions and claims that shaped the philosophical landscape as surveyed in Broad's 'elementary reflections.' As Mike Martin (2006) and others have argued, it makes at least as much sense to see some representationalist accounts of perception's phenomenal character as inheriting one half of the sense data view's divided legacy, with another view, 'naïve realism', receiving the other half.

We can get a sense of this way of seeing the contemporary discussion by considering the way Broad organizes his own. He begins by considering perceptual experience from a "purely phenomenal point of view" (pp. 3-4). This is supposed to be grounded in the way things appear, both to "unsophisticated percipient[s]" and "even to sophisticated percipients whose knowledge of the physical and physiological processes involved assures them that the appearances are largely misleading" (p.4). Naïve realists (then and now) take the side of the 'unsophisticated percipient' and contest the judgment of her sophisticated foil. They think we should believe the appearances; and, according to Broad, the appearances are such that in perception we are aware of (we "prehend"), among other things, objects in the environment and their properties, as well as processes and events with either visual (flashes of lightning, pulses of illumination) or auditory (the tolling of the bell, the roar of a waterfall) qualities. Perception is a mode of relation to these things – the terms that has stuck is 'acquaintance' – and is as it is, phenomenally, at least partly because of the way these things are.

In the terms I'll use throughout my discussion, this is an answer to a question about the 'scope' of perception grounded in a claim about its metaphysics. Perceptual scope questions are questions about what subjects perceive. A common way of asking this kind of question is to ask what the contents of perception are. But the question is intelligible even if, as some people think, perception has no content. Even if perception is contentful, however, the question Broad is interested in involves more than content since it is a question about how the putative objects or content of experience appear, and this involves more than can be specified by specifying a

content – minimally it involves the way that content evolves over time.[1] It might also involve the way that different contents (or their non-contentual analogues) fit together in an experience. We can call these additional aspects of a perceptual experience its structural features. Assuming these make a difference to the appearances, structural features of perceptual experience are relevant to the question of what falls within the scope of perception. If perception is grounded in a relation, then this relation can be appealed to in answering the scope question – and this involves saying both what a subject is perceptually related to and how. Naïve realism is in this sense a purely structural approach to the scope question. It attempts to answer it without talking about content – or by only taking content to be derived from the some more basic psychological relation, realized in perception.

So, the unsophisticated, 'naïve' answer to the scope question cites objects, events and their properties in the environment to which subjects are related. Since these are part of the 'physiological and physical processes' involved in perceiving a (*genuinely* naïve) member of Broad's audience might wonder what dislodged the sophisticated percipient from these natural-seeming conclusions about the scope of perception. But someone familiar with either contemporary or contemporaneous debates in the philosophy of perception could likely guess the answer to this question. Introduced in very general terms: it is a necessary condition on perceiving anything that in addition to the facts that seem apparent from the purely phenomenological point of view "a certain internal change is transmitted from the stimulated sense-organ to a certain part of his brain and sets up some kind of disturbance there" (p. 10). The source of the sophisticated subject's assurance that things aren't really as they seem stems from this. The reason is that although "[t]he physiological and anatomical facts just stated do not suffice to prove that seeing, hearing, and feeling are not […], prehensions [i.e., direct perceptions] of events and of certain of their intrinsic qualities", they do suggest a critical problem with the naïve view:

> What would be fatal to the [naïve] account of these experiences would be if it could be shown that if the area of one's brain were suitably affected, such experiences might occur even though no such object were then occupying the space in question.

This is the set up shared by the mid-century and the contemporary discussions. Philosophers continue to provide arguments aimed at establishing or resisting this negative conditional claim. A well-worn approach focuses on criteria by which the sameness of an experience in the normal

[1] The simplest way to think about this question is as a question of how a series of contentful states unfolds. But stating this involves relations between states not given in their content.

case with the experience one might have even if there were no object one perceived might be determined. Defenders of the unsophisticated view provide increasingly sophisticated accounts of the object-involving conditions an experience would have to meet, over and above the brain's being 'suitably affected' by occurrent internal conditions – certain internal changes disturbing certain internal parts.

But an alternative approach has been received less attention from naïve realists. Broad (who falls on the other side of the debate) touches briefly on one element of this alternative but rejects it out of hand. We do not, he states, distinguish between waking experiences that are perceptions and dreams that are not "by noting dissimilarities in their phenomenological character" rather, "[w]e do so by considering *the inter-relations* of experiences with the earlier and later experiences of the same person" (p. 12, emphasis in original). Here Broad is probably paraphrasing Descartes' familiar point that the ordinary introspective criteria for distinguishing dreams from episodes of waking consciousness are provided by relations between successive experiences. If so, then the phenomenal character he is considering is revealed by introspection and temporally bounded in a way that allows us to make sense of the idea of relating and comparing distinct experienced moments. Since differences in phenomenal character are contrasted with the differences in the way experiences are related, this suggests that each moment has its phenomenal character independently of the others: at one moment you might be carrying on a conversation with a friend, at the next you might be scuba diving – you can tell you are dreaming not by reflecting on the character of either of these 'experiences' but on the way they fail to cohere into a plausible experiential episode.

Leaving aside some substantive and questionable assumptions made here about the phenomenology of dreaming (for discission see Dennett 1978, ch. 8), we might also challenge the opposition that Broad makes between the phenomenal character of an experience and its interrelation with other experiences. Many philosophers think of experience as episodic and as constituted by their place within the wider episodes that make up the part of people's psychological lives of which they are (perceptually) aware (see, e.g., Steward 1997; Martin 2002; Soteriou 2013). Does it, on the other hand, make sense to think of phenomenal character atemporally? What about the processes that underlie it? This second question is at the heart of the 'enactivist' movement in cognitive science and philosophy (Ward et al. 2017). Enactivists are motivated by the thought that perception and action are processual, and in particular that they constitutively involve an agent's active integration with – or even 'enaction' of (Valera et al. 1991) – her environment. While similarities between enactivist and naïve realist approaches have been noted, few philosophers until recently have laid out and made good on these

parallels.[2] There are several reasons for this, but one important one is related to the contrast noted above between the current and the earlier debate about sense data and representations. The though is that enactivists and naïve realists pitch their accounts at different 'levels of explanation', with naïve realists opposing 'intentionalist' theories at the personal level of description with enactivists opposing later, 'subpersonal' representationalism. From one perspective, naïve realists continue to rehash the hackneyed concerns that animated Broad's reflections on Pritchard and Russell or Austin's critique of Ayer; from the other, enactivists engage in armchair speculation on essentially empirical disputes (which seem, in any case to be decided against them).

I think this is unfortunate, and this thesis aims to contribute to an emerging effort to upset this settled dynamic. Like some other participants in this discussion (see Raleigh 2021; Gomes 2016, p. 946; and, especially, Ward 2016; 2023), I take the work of Susan Hurley (1998; 2003a; 2008) as a guide to how this might be done. Hurley's contributions to philosophy of mind fall in a tradition shaped by both Kantian and ordinary language critiques of the empiricist outlook from which sense data views developed, and it articulates a model of the relation between personal and subpersonal levels of description that rejects its contemporary representationalist analogue. It does this while endorsing and incorporating mainstream cognitive scientific methodology and its ubiquitous appeals to representation – all while defending an answer to the scope question that, like the naive account, adverts to the subject's relation to environmental objects. This account inspires what I'll call the sensorimotor acquaintance view. Outlining the framework into which this account falls, I'll make the case that people motivated by the personal-level considerations behind naïve realism have good reason to adopt something like Hurley's model of the personal/subpersonal distinction.

b) Methodology and assumptions

The argument of this thesis is shaped by a methodology set out across Hurley's work, whose development and discussion of Dennett's (1969) distinction between personal and subpersonal levels of explanation traces another connection between the earlier debate about sense perception and the contemporary debate about phenomenal character. Despite the starkly different conclusions reached by Hurley and Dennett, I understand them both as formulating versions of what I'll call an interpretationist model of perceptual content. Interpretationists think that our account of how psychological phenomena have the contents we ascribe to them

[2] Some resent attempts to do this systematically include Raleigh (2021), Ward (2023); for earlier discussion, see Campbell (2008), Martin (2008) Noë (2008)

should start from reflection on what goes on in the ordinary practice of psychological interpretation and explanations of people as more-or-less rational agents, motivated by and responsive to their standing attitudes and perceptual experience.

A ubiquitous feature of these practices is the attribution of content that characterises persons as being in states, undergoing episodes, that in combination with other attributed contents makes sense of their speech, behaviour and action – that makes them intelligible as subjects of experience and responsible agents. Intelligibility imposes constraints on behaviour and relies on context for interpretation, and it is against this backdrop that the explanatory practices that support content attribution emerge – in a slogan: context determines content. While the interpretationist framework is appealed to in support of a strict demarcation between personal and subpersonal levels, Dennett and Hurley suggest, in different ways, a more integrative approach.

I aim to use an interpretationist approach along these lines to support claims about the phenomenal character of perception. These are scope claims about a class of perceptual experiences I'll call perspectival experiences. The next section of this introduction will begin to set out the idea of a perceptual perspective that unifies this class. But first, I'll say some more about how I will be understanding the notions of *phenomenal character*, *representational content* and *interpretation* and their interrelation:

> *Phenomenal Character*: Discussions of phenomenal character often introduce it in first personal terms – for a state to have a phenomenal character there is for there to be *something it is like* to be in that state (Nagel 1974). (Likewise, for an episode to have phenomenal character is for there to be something it is like to undergo the episode). Implicit in this way of introducing phenomenal character is an appeal to introspection – since we know in our own case what it is like to X, we have our own stock of examples from which to build towards a more general characterisation: introspection on experience gives access to types of appearances, which we can assume are tokened by other people when we attribute experiences of this type to them. Doubts about this intuitive picture motivate a variety of alternatives – providing objective criteria for 'appearances', replaying introspection with 'heterophenomenology', locating phenomenal character in objects of experience rather than in introspectable psychological episodes considered independently. The approach this thesis takes involves aspects of these mentioned approaches: aiming to characterise an aspect of phenomenal character of some episodes of experience as the phenomenal contribution

made by the objects of those experience. Here the notion of content is helpful, since it captures an aspect shared by any attempt to articulate phenomenal character.

Representation: The notion of representational content is ubiquitous in cognitive science and controversies about 'representationalism' are nearly ubiquitous in philosophy of mind. It is not always clear how much the notion used by cognitive scientists and the notion argued about by philosophers overlap. In this thesis, I'll suggest that the way in which way of talking about representation, allows us to sidestep much of the controversies that animate the philosophical discussion while allowing the explanatory use that it is put to in cognitive science to inform the debate about the scope of perspectival experience. On this view, content tracks the specific contributions made to an episode of experience by that experience's object – that is, the contents attributed to perceiving subjects when they have a perspective on an object have singular content specifying that object. But they also identifying commonalities across experiences and across perceivers – properties an object is represented as having are part of the representation's general content. Both of these features are part of the cognitive scientific notion, but they also give us what we need for the questions dealt with by this thesis.

Interpretation: Crucial to the explanatory payoff advertised above is an understanding of content attribution in cognitive science as tied up with practices of content attribution that are equally ubiquitous in everyday life. 'Folk psychology' (Dewhurst 2021) is more than 'propositional attitude' attribution, but it has at its core the interpretation of others in terms of the contents attached to states, episodes and events that they occupy, undergo and experience. Belief and desire are paradigms, but perceptual and affective phenomena are – in ordinary explanatory practice at least – intentional. They have contents given by specifying the actual or unactualized things they are about. Where content attribution 'goes subpersonal' (Hurley 1998), it relies on the criteria of interpretability established at the 'personal level'. This makes sense of the mechanisms and processes that underly content individuated phenomena on a model established at the personal level: interpretability in terms of contents is a guide to discovery of mechanisms, that make sense of the subjects as an embodied agent in causal interchange with their environments.

My goal isn't to defend the commitments taken on in these summaries. Each of them by itself contains more than enough material for another thesis. Instead, I will try to illustrate

connections between the ideas at the centre of each and suggest that they can be used to make sense of the main, attractive set of claims held in common by contemporary naïve realists and enactivists. I'll suggest that if we understand contents as interpretations and interpretations as dependent on contexts, then some of the main claims that both groups want to make about the constitutive role of the environment in at least some perceptual experiences fall naturally out of the cognitive scientific picture.

That it hasn't often been seen this way itself needs to be explained. Providing an explanation is the second way this thesis aims to contribute to the contemporary debate – by drawing attention to a second distinction made by Hurley between types of *subpersonal* explanation. Some explanations in cognitive science identify correspondences between mechanisms and personal level phenomena that cross the two levels. For example, subcortical mechanisms are said to compute 'rewards' (Ross et al. 2012) or early visual areas to provide a map of the subject's 'visual field'. These accounts rely on what Hurley (1998) calls 'interlevel isomorphisms' if contents are shared between levels or if properties of phenomena at one level (e.g., temporal dynamics, causal powers) are projected across levels. Hurley's version of sensorimotor enactivism is grounded in a critique of explanatory isomorphisms in the case of perception and action and her development of a non-isomorphic model of interlevel explanation.

The notion of explanatory isomorphism, central to *Consciousness in Action* and Hurley's later work in the philosophy of cognitive science, needs some introduction. The relevant isomorphisms are between models used to characterise phenomena at respectively personal and subpersonal levels of explanation. An isomorphism holds where elements of one model can be mapped one-to-one onto elements of the other. This mapping can be reversed, so that either of two isomorphic models can be used to predict (given the mapping function) every element of the other model. Consider a simple subpersonal model of perception that specifies input states of the sensory array and internal states that are elicited by sensory input. An explanatory isomorphism holds if a psychological model ascribes perceptual states to a person in a way that mirrors the corresponding subpersonal model so that its sensory input and elicited states always correspond to the personal level perceptual input and resulting perceptual states. This kind of strict isomorphism at the levels of models used to characterise phenomena across levels is (like the notion of levels itself) an idealisation. This ideal character is what allows isomorphisms to hold – *between models.* Meanwhile, an *explanation* is isomorphic if it goes by appeal to modelled interlevel isomorphisms of this kind; if it does not, it isn't.

Hurley's distinction between isomorphic and non-isomorphic explanations offers a helpful perspective on many debates in philosophy and cognitive science, of which the discussion around perceptual acquaintance is one. Explanatory isomorphisms between neural mechanisms and perceptual states exclude the kind of structural features required for acquaintance since acquaintance-theoretic models of personal-level perception contain elements (perceived objects, environmental causal relations between these and perceivers) that are excluded from subpersonal models of mechanisms. But while Hurley rejected explanator isomorphisms in the cases she focused on, her model of interlevel explanation does not rule them out generally. Approaching debates about psychological phenomena with prior commitments either for or against explanatory isomorphism is a mistake since it encourages to read substantive metaphysical claims off the essentially methodological distinction between explanatory levels. While I'll argue that in the case of psychological phenomena metaphysics should track explanatory methodology, it's a consequence of Hurley's model that it does so on a case-by-case basis.

c) Acquaintance, alienation and the subject of perspectival experience

Hurley is often presented as a pioneer of sensorimotor enactivism in the philosophy of cognitive science – with her critique of isomorphic explanation cited as an ongoing inspiration by some contemporary enactivists (see, e.g. Kirchoff & Kiverstein 2019). But her work develops out of a tradition of 'Oxford Realism' from which contemporary naïve realism also emerged. Hurley's work thus marks a point where the roots of enactivism and naïve realism intertwine; and central to this underlying commonality is the notion of a subject's perspective on objects in her environment. Naïve realists and enactivists deny that perspective can be accounted for in representational terms. Rather, answers to the scope question about perspectival experience appeal wholly or in part to its structural features.

The view I find in Hurley is that perspective involves both structure and content. While she doesn't use the term 'acquaintance', her notion of 'perspectival experience' captures a similar set of structural features of the subject's relation to environmental objects – with the crucial addition (or emphasis) that perspective is something that the subjects, as embodied agents, actively bring about. For her, the notion of non-isomorphic explanation makes space for perspectival experience because having a perspective involves interdependencies between perception and action that cannot be fully specified in terms of representational mechanisms in the brain. These dependencies run through the environment and the perceiver-agent's body and so locate her at the centre of a causal network that determines, or so I will argue, the

perspectival structure of her ongoing experience. What makes someone interpretable as having a perceptual perspective is her location within a causal network of this kind – something that is revealed by and depended on by psychological explanation at both levels. Isomorphic explanation of a person's perceptual states in terms of occurrent representations in perceptual mechanism truncate this underlying causal structure and so give rise to a picture of perception that denies that objects play a constitutive role in experience, further giving rise to the problems that both naïve realists and enactivists respond to.

Common to Hurley, contemporary enactivism and naïve realism is a concern to avoid a distinctive kind of misconception about the subject's place in the world that I will call 'alienation'. The notion of alienation is a slippery one, but it shows up repeatedly in enactivist and naïve realist critiques of mainstream views in philosophy of perception and action. Voicing a version of this thought, Hurley (1998, p. 36) identifies cognitive science as a potential source of alienation:

> Increasing knowledge of the brain can make it hard to accept what must once have been taken for granted: that what it is to be a subject and an agent is open to view just where it seems to be, in the lives of us familiar middle-sized objects, persons, as we interact with our normal environments.

Picking up on this suggestion, we can understand alienation as a disruption of the picture of the agent and subject which Hurley describes as having been taken for granted. Of course, making sense of this negative characterisation involves the positive task of setting out the unalienated picture. This is what naïve realists' evocation of the ordinary language we use to talk about perception is meant to provide. If what it is to be a subject and agent is 'open to view', then each of us has access to it – at least in principle. This easy access to the unalienated picture is a central theme of naïve realist approaches to perceptual experience, which have been grounded both in introspection and in attention to the folk-psychological language of ordinary interpersonal explanation and interpretation. This thesis takes the second of these approaches, emphasizing the role of folk psychological interpretation at the personal level in setting out the normative and causal structure of the phenomenon to be explained: the perspectival unity of the perceptual experiences that subjects like us enjoy.

Bringing the subpersonal level of explanation into the interpretive mix has the potential to refine this picture by revealing more about the nature of the agent-environment interactions that underly perspectival experience. As all its above-mentioned critics agree, the alienated picture arises from taking a certain view of this underlying structure – as Broad's 'internal changes' and

'disturbances'. But it remains unclear just how much of what naïve realists find in our ordinary talk about perception can be retained. Enactivists propose that the subpersonal realisers of perceptual experience include patterns of engaged interaction between perceivers and their environment. For *sensorimotor* enactivists, who draw most on Hurley's work, these patterns are best understood as causal loops mediated by sensory and motor systems. Marrying these approaches, then, has the potential to cast light on – but also to limit the uses of – a relation of perceptual acquaintance enabled by sensorimotor interactions. This is the goal of this thesis, which draws on these overlapping literatures to put forward a *sensorimotor acquaintance view* of perceptual experience.

The sensorimotor acquaintance view understands ordinary perception as realising a concrete relation between subjects and environmental particulars that constitutes these particulars as perceptual objects and enables a subject's thought about and action on them. This relation is grounded in sensorimotor processes that unfold over the period that a perceiving subject and an environmental particular are so related. It explains the phenomenal character of the experience as a function of properties of both relata; and so, on this view, perceptual objects play a constitutive role in the phenomenology of perceptual episodes. This constitutive role performs some but not all of the explanatory roles that many contemporary philosophers assign to the notion of perceptual acquaintance. Understanding the subpersonal and personal-level structure of perspectival experience, we recognise the essential contribution made by the objects that make up the environment that embodied perceivers bring into view and have – in virtue of their embodied nature – a perspective on.

d) Thesis outline

The thesis is divided into five long chapters, grouped together into two parts. Part I assembles the components out of which the sensorimotor acquaintance view is made up; Part II fits them together. Each chapter explores one or more subplots, and some of the key ideas set out in this introduction – for example, the contrast between alienation and acquaintance – recur across multiple chapters in different contexts. To help keep track of these, each chapter introduction includes an overview like this one, identifying further themes and giving more detail on the chapter's structure than is provided in this short summary.

Part I is made up of three chapters that focus on different components of the sensorimotor view, giving exposition and a sense of the philosophical problems that makes appealing to them necessary. The main components are the notions of acquaintance (§1), subpersonal explanation (§2) and perspective (§3). These chapters attempt to give a historically informed introduction to

the terms in which the positive argument of Part II is made, but they do not give a comprehensive survey. Instead, I've aimed to provide a narrative based on the work of some key thinkers, which supports the characterisations I give and go on to use.

The conclusion of the positive argument given in Part II is that perspectival experience is fundamentally relational: its structure is determined by a concrete relation between a subject and objects, and this structure matters to how philosophical questions about perception should be answered. My presentation of this argument is divided into two chapters, corresponding to the relata of the sensorimotor acquaintance relation: environmental objects (§4) and perceiving and acting subjects (§5). Chapter (§4) gives an argument that once assumptions in favour of explanatory isomorphisms for perceptual experience are abandoned, the objects of perceptual experience are revealed to be the traditional objects of acquaintance: ordinary environmental objects. Chapter (§5) argues that given the methodological claims argued for above, we should appeal to this relation to answer scope questions about perspectival experience. It concludes that one central naïve realist claim – that environmental objects play a constitutive role in our experience of them – is correct.

Part I

Perceptual Acquaintance and Levels of Explanation

1

Naïve Realism, Neurocomputation and Perceptual Acquaintance

0) Introduction

This chapter lays the ground for my defence of the sensorimotor acquaintance view in two ways. First, it loosely sketches the history of 'acquaintance' as a technical term in philosophy of perception. I outline how anti-idealist and anti-sceptical programmes in analytic philosophy, instigated at Cambridge and Oxford near the beginning of the last century continue to inform the use contemporary philosophers make of the notion of perceptual acquaintance. This is a tradition that Hurley's work fits into; but it also, from its beginning, has been motivated by strong 'realist' intuitions that might seem to be in tension with the neo-Kantian elements of her project. This chapter makes a start at working out how and at what points these apparent tensions might be relaxed.

Second, the chapter differentiates the sensorimotor acquaintance view from another way of thinking about perceptual experience that also draws on the legacy of Oxford and Cambridge realism. This is articulated in the contemporary revival of 'naïve realist' views (see, e.g. Martin 2002; Campbell 2002; for discussion, Genone 2016; Raleigh 2019; French & Philips 2023) In particular, I will focus on and criticise a recent attempt by Beck (2019) to articulate a version of naïve realism that explicitly connects its picture of the perceiving subject to neurocomputational models of its implementation. Where this account goes wrong – or so I'll claim – is that it tacitly relies on the assumptions about how subpersonal explanation of perception works, which I'll discuss as explanatory isomorphism in chapter (§2). These turn out to be incompatible with the naïve realist commitments that Beck takes on. But, as I'll argue in the later chapters, the most promising non-isomorphic alternative model delivers more direct support to the sensorimotor acquaintance view than to naïve realism.

Here's the plan for the chapter. Section (§1.1) fills in some of the historical background to the debate about naïve realism, focusing of the connections between acquaintance and claims about the 'transparency' of perceptual experience. Section (§1.2) maps out a part of the contemporary philosophical landscape that these views fit into and begins to motivate some

doubts about the naïve realist picture. Section (§1.3) outlines how Beck (2019) responds to these. This involves explicitly connecting naïve realist claims about the phenomenal character of perceptual experience to neurocomputational accounts of its subpersonal implementation. I argue that Beck's proposal faces serious problems, which the following chapter's discussion of isomorphic subpersonal explanation will put us in a position to appreciate.

1) Acquaintance, Transparency and Realism about the Objects of Perceptual Experience

a) Acquaintance and the rejection of idealism

The claim that a distinctively perceptual acquaintance relation allows environmental particulars to determine phenomenal properties of experience is nowadays frequently associated with a view in the philosophy of perception called 'naïve realism'. For example, Soteriou (2009) describes the core naïve realist commitment as the claim that:

> perception is an acquaintance with mind-independent empirical things whose basic natures and perceptible qualities constitute what it is like to be presented with them in this way.

French and Philips, meanwhile, characterise naïve realists as holding that:

> perceptual experience involves a subject standing in a primitive relation of perceptual acquaintance to aspects of mind-independent reality which constitutively shape the contours of the subject's consciousness (2023, p. 2)

In the last twenty years, appeals to 'acquaintance' to explain fundamental epistemic, semantic and phenomenal features and capacities of the mind have undergone a "renaissance" (Raleigh 2019), and naïve realism can be understood as the perceptual wing of this wider movement (for further recent discussion, see Genone 2016; Siegel 2019; Duncan 2021).[3] The way this story is told continues to shape these debates, so it will be helpful to get clear on some of its relevant details.

The main forerunners to this notion of acquaintance play an important role in what Gomes (2017, p. 529) calls the "textbook nativity story" of analytic philosophy. One thing to draw from this is that many original formulations of naïve realist ideas were grounded in a rejection of

[3] Here, I'm following Gomes (2017) in marking 2002, in which Martin's 'The Transparency of Experience' and Campbell's *Reference and Consciousness* were published, as a watershed year for contemporary naïve realism. While my discussion will touch on acquaintance-based proposals that were advanced before it, notably those developing Evans' (1982) Russellian account of singular reference (e.g. McDowell 1994), I'll mostly mean by 'naïve realism' the view that was definitively on the table after this point.

(allegedly) Kantian models of the relation between subject and world (pre-eminently, e.g., Pritchard 1909). While some contemporary versions explicitly maintain this commitment (Campbell 2011; Kalderon 2011), other commentators have meanwhile found naïve realist elements in the Kantian picture (Gomes 2014, 2017; Allen 2020); and McDowell's (1994) formulation of perceptual disjunctivism is self-consciously a Kantian account. Since the positive proposal I make in the next part of the thesis ties claims about perspectival unity developed within the 'Kantian' branch of this broad realist tradition to an account of the objective structure of the perceptual relation grounded in acquaintance, another reason to outline this history is to work out how far this project can share in the general motivations behind the renaissance of acquaintance.

Probably the most familiar part of perceptual acquaintance's origin story – the part shared with the conventional nativity of analytic philosophy – revolves around on Cambridge and the work of Moore and Russell. This was the context for Russell's 'Knowledge by Acquaintance and Knowledge by Description' (1911), in which 'acquaintance' is, arguably,[4] first used in roughly the sense that it has for contemporary philosophers of perception. A simple version of the story goes like this: Moore and Russell reject a neo-Hegelian idealist picture of the relation between subject and world, which had shaped philosophical orthodoxy in Cambridge in the late 19th Century. Developing Frege's insights about the semantics of referring expressions, Russell sets out an ontology of discrete particulars and universals, with the important epistemological consequence that objects of knowledge are independent of psychological acts or processes involved in knowing them. Next, Moore (1903) argues that this ontology entails a 'refutation' of idealism, since in knowing about the objective world, subjects come to be related to entities distinct from themselves. Later, Russell (1911, 1912) introduces the term 'acquaintance' for the most basic and direct of such relations:

> I say that I am *acquainted* with an object when I have a direct cognitive relation to that object, that is, when I am directly aware of the object itself. (1911, p. 108)
>
> The faculty of being acquainted with things other than itself is the main characteristic of a mind. Acquaintance with objects essentially consists in a relation

[4] Raleigh (2019) points out that Russell uses the term in a recognisably similar way at the beginning of 'On Denoting' (1905) and mentions several precursors, who use 'acquaintance' (James 1890) or 'das Kennen' (Helmholtz 1868) to refer to a non-propositional, direct mode of knowing, paradigmatically manifested in sensory perception. Meanwhile, apart from the word itself, another recent review of the philosophical literature claims that acquaintance "or something very much like it shows up throughout the history of Western Philosophy – going back at least to Plato" (Duncan 2021). Nonetheless, the scope of the current discussion is (mainly) limited to the 20th Century sources of contemporary naïve realism.

> between the mind and something other than the mind; it is this that constitutes the mind's power of knowing things. (1912)

This account, in which 'acquaintance' fills something like the role given by later philosophers and cognitive scientists to 'representation' as the general means by which the mind comes into contact with the world, provides a model for contemporary naïve realists. It's worth noting, then, that 'cognitive' in the first quotation doesn't have its broad, post-cognitive revolution meaning; it refers specifically to a kind of knowing. The acquaintance relation doesn't just explain knowledge, it is the most fundamental kind of knowledge. But, importantly, knowledge by acquaintance is not "the sort of relation which constitutes judgement, but the sort which constitutes presentation" (Russell 1911, p. 108). Judgment relates subjects to propositions, which are structured composites of particular and universal elements (cf. Stevens 2008). Presentation involves a simpler relation between a subject and an object of acquaintance. This is the subject's most basic way of being related to mind-independent reality.

While a subject's being acquainted with some object, o, doesn't by itself constitute a judgment *that* o has some property, it does play an essential role in her being able to make such judgments. To be presented with an object is to be directly aware of its the properties. In combination with her other cognitive (and possibly linguistic) capacities, this is then a basis for the subject to judge that the object has those properties (on Russell's developed view: to stand in the judging relation to a propositional complex that includes the object and property as elements). For Russell and Moore, perceptual acquaintance was an element in a more general 'direct realist' project (Hylton 2003; Griffin 2012). While Russell (1912) discusses perceptual acquaintance, he also claims that subjects can be acquainted with universals, their own memories, and (possibly) with their 'selves' through introspection. These are constituents of further epistemic abilities, which come together in the exercise of complex capacities such as judgment. Perception is essentially one expression of this basic capacity being presented with elements of mind-independent reality. It is knowledge by acquaintance of the immediate objects of sensory perception and their properties.

Russell and Moore's turn away from idealism is sometimes thought of as a return to British empiricist models of the epistemology and metaphysics of perception. This impression is encouraged by the way they, and many other early analytics came to understand the objects of perceptual acquaintance. By the time Russell had produced the characterisations of acquaintance quoted above, both he and Moore (Moore & Stout 1913; see also Snowdon 2015) accepted sense-data theories of perception. On this kind of view, perceivers are acquainted with

basic sensory qualities (e.g. heat, tone, colour, etc.), which are the properties of their sense data. Ordinary material objects and other features of the perceptible environment stand in relations of "correspondence" with sense data and are known 'indirectly' via acquaintance with these immediate objects.[5]

Nonetheless, we should be careful not to underestimate how radically Moore and Russell's project aimed to break with previous approaches. In a frequently quoted passage, Russell (1959) narrates his entry into the 'new philosophy' as a departure not only from the contemporary British idealism, but from a longer tradition stretching back through "Kant and Hegel" (p. 42) to early modern conceptions of perception as constituted by affectations of or impressions on the mind of the subject:

> I felt it, in fact, as a great liberation, as if I had escaped from a hot-house on to a wind-swept headland. I hated the stuffiness involved in supposing space and time were only in my mind. I liked the starry heavens even better than the moral law, and could not bear Kant's view that the one I liked best was only a subjective figment. In the first exuberance of liberation, I became a naive realist and rejoiced in the thought that grass is really green, in spite of the adverse opinion of all philosophers from Locke onwards. (p. 48).

This account, in which naïve realism figures as a stage in the development of Russell's mature view, suggests that the problem for which the Cambridge realists' relational theories of knowledge and perception were supposed to provide the solution had roots going deeper than the influence of Hegel – or of any particular contemporary account. For Kant, space, time and the empirical objects that populate them are (as we experience them) all products of the transcendental activity of the subject in coming to know them. For Locke and the empiricists, the greenness of the grass is a feature of sensory impressions that are internal to the mind. Understanding the mind as fundamentally relational was a way to avoid making any of these subjectivist commitments, which had shaped the philosophical presuppositions of the thinkers that the early analytics were reacting against. While abandoning naïve realism, Russell insists

[5] Moore's view at this point was closer to the naïve picture than Russell's. For example, he suggested that visual sense data couldbe identified with the facing surfaces of objects (Moore & Stout 1913), while Russell insisted that the perceptible properties of sense data could not be properties of ordinary physical objects or their parts. However, both agreed that perception of whole physical objects depended on inference from sense data. Later in his career, Moore came to accept a view much closer to Russell's 'Berkelian' account (for discussion, see Snowdon 2008, Travis & Kalderon 2013).

that he retained the insights that had freed him from the "subjective prison" (ibid.) constructed by the modern philosophical tradition.

We can get a sense of the Victorian hothouse Russell saw himself to be escaping from by comparing Russell's recollections of his exuberant early naïve realism with Walter Pater's description of the decidedly un-naïve perception afforded, as he understood it, by philosophy and aesthetic criticism:

> At first sight experience seems to bury us under a flood of external objects, pressing upon us with a sharp and importunate reality, calling us out of ourselves in a thousand forms of action. But when reflexion begins to act upon those objects they are dissipated under its influence; the cohesive force seems suspended like a trick of magic; each object is loosed into a group of impressions—colour, odour, texture—in the mind of the observer. (1873 [2001], p. 208-209)

Strikingly, this is a picture of the perceptual situation is presented as a discovery – it is not just a reflective stance one could take to experience but reveals something about its fundamental nature. Understood in this way, Pater's description captures something that the views that Russell and Moore set themselves against have in common.[6] What should stand out, in this context, is a characteristic reversal of appearance and reality. As Moore (1903) complains of the "modern idealism", a consequence of accepting this kind of picture is believing "that the universe is very different indeed from what it seems" (p. 433). While experience *seems* to present us with the objects of everyday life, calling us out of ourselves and eliciting action, philosophical analysis reveals their solidity and cohesion to be illusory. What there *really* is, meanwhile, are groups of essentially mind-dependent impressions, "which burn and are extinguished with our consciousness of them" (Pater 1873,p. 209). From this point, it's only a short step to a kind of subjective idealism:

> [e]xperience […] is ringed round for each one of us by that thick wall of personality through which no real voice has ever pierced on its way to us, or from us to that which we can only conjecture to be without. Every one of those impressions is the impression of the individual in his isolation, each mind keeping as a solitary prisoner its own dream of a world. (ibid.)

[6] See Inman (1981) for discussion of the intellectual context of Pater's conclusion to *The Renaissance*, which sets out an eclectic mix of influences, including 19th Century physiology, British empiricists like Hume and Locke, Hobbesian materialism, and the German idealists Kant, Fichte and Hegel.

For Russell and Moore, passages like these are symptomatic attempts to make sense of the mind and its contents independently of its relations to mind-independent reality. The mind, for the Cambridge realists, is a system of relations; and its perceptual part consists of the relations that tie it to empirical reality. This is where analysing perception in non-relational terms gets us. If perception does not, constitutively, put subjects in touch with a world that is independent of and external to them, then there is no sense in which the ordinary objects of perception can be understood apart from their appearance on the stage inward stage of a purely subjective 'Cartesian theatre' (Dennett 1991).

So, while acquaintance is introduced as an epistemic relation – a mode of knowing the mind-independent world – the work it is put to by Russell and Moore involves claims about the metaphysics of the perceiving subject's relation to mind-independent reality that are tied up with their particular anti-idealist project. Three of these have been particularly important to contemporary naïve realists. These characterise perceptual acquaintance as:

> *Basic*: There is no more fundamental mental state or activity that could exclude the object while playing the same role as acquaintance in the psychological life of the subject.
>
> *Transparent*$_{Ext:}$ Perceiving subject are aware of (properties of) objects of acquaintance. The perceptible properties an object has independently of the acquaintance relation are the properties that the subject is acquainted with.
>
> *Transparent*$_{Ins}$: In introspection on perception, subjects are aware of the (properties of) objects of perceptual acquaintance.

Considered in isolation from the wider metaphysical commitments of the Cambridge realists (instead, e.g., thinking of 'acquaintance' as roughly equivalent to a pre-theoretical sense of 'awareness'), it's not obvious that these claims should go together. Only by combining a sense data theory of the objects of perception with the anti-idealist project outlined above did the Cambridge realists come to see these claims as bound up with each other – and to propose strong versions of each.

A metaphor that is sometimes used to introduce these sense data views (particularly for visual perception) is the notion of a 'property mosaic' made up of sense data that fills the subject's phenomenal field. They claim that every aspect of a subject's experience is determined by the properties of sense data and their arrangement (this is what *transparency*$_{Ext}$ amounts to); and when a subject introspects on her experience, it is just this mosaic, and its arrangement that she

is presented with (*transparency*$_{Ins}$).[7] Moreover, there is no other way for a subject to be in a perceptual state with a particular phenomenal character than for an arrangement of sense data with the appropriate properties to be presented to her (*basicness*). So, for example, veridical perceptions and subjectively indistinguishable hallucinations are both presentations to the subjects of equivalent sense data mosaics.

Questions about the extent to which these claims should be modified when the objects of acquaintance become ordinary material things are at the core of the contemporary debate about naïve realism. In particular, this concerns the viability of various strategies for responding to arguments from hallucination. Before discussing these, however, it will be useful to get the other half of naïve realism's conventional origin story into the picture.

b) Oxford realism and perceptual knowledge

At roughly the same time as Moore and Russell were developing their anti-idealist programme in Cambridge, a corresponding realist tendency was developing in Oxford. Marrion (2000) and Travis and Kalderon (2013) identify its primary figures as Cook Wilson and Pritchard. These laid the groundwork for a distinctive 'Oxford view' of language, epistemology and perception that came to be most completely expressed later, by Austin. The perception wing of Oxford realism remains influential and, as we'll see, supplies a missing epistemological ingredient required for contemporary naive realism to develop.

Discussing the above-quoted passage from Russell (1959), Travis and Kalderon remark that, since both he and Moore endorsed sense data theories of perception, his claim to have retained their early anti-idealist insights can only be "half right", since neither "was able to hang onto that realism from which they began" (2013, p. 491). But if 'realism' here is supposed to contrast with the 'idealism' or 'subjectivism' discussed above, then Russell and Moore could both fairly

[7] This picture is already complicated in Moore's (1903) canonical statement of the introspective transparency thesis. Famously, he claims that "the moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to introspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous." That is, sensation is diaphanous to what it is a sensation of, which is a property of its object (i.e. in this case, blueness). But *diaphaneity* is not perfect transparency, and Moore goes on to suggest that the sensation can itself be an object of introspection. This allows it to be a distinguishable element in introspective experience. In fact, he identifies the main task of the passage in which the transparency/diaphaneity claim is made as getting his reader to appreciate the phenomenal presence of this separable subjective element (ibid.). If Moore's (1904) objects of perceptual awareness are assimilated to (mind-independent) sense data, it looks like this proposal would rule out the strong reading of transparency$_{Ins}$ just outlined. As Travis and Kalderon (2013) note, however, Moore's later use of the term sense data was idiosyncratic – and it was only towards the end of his career that Moore adopted the more familiar Russellian picture. Russell's (1912) view plausibly does involve the strong claim.

contest this assessment. To see why, think about the differences, highlighted in the previous section between sense data as conceived by the early analytics and the classical empiricist notion of sensory impressions.

Russellian sense data, unlike sensory impressions, are mind-independent particulars – they are not modifications of, or events within, a perceiving subject's mind, but are rather those parts of the objective world most immediately available to consciousness.[8] Understanding perception as acquaintance with sense data, so-conceived, is supposed to represent an improvement on the hothouse subjectivism exemplified by Pater (1873) just because they are continuous with the rest of the world apart from the subject. The thought is that perceiving subjects can make sense of a world that is constitutively distinct from their experiences of it, since they are already acquainted with things (sense data) that are as they are independently of the psychological acts or processes involved in becoming so acquainted.

Nonetheless, this anti-idealist proposal is compatible with believing, as Russell later puts it, that "common sense leaves us completely in the dark about the true intrinsic natures of physical objects" (1946. His and Moore's (1903) complaint against idealism was not that it suggests, generically, that the physical universe is other than it seems, but that it makes a specific implausible claim about what this difference amounts to:

> Chairs and tables and mountains seem to be very different from us; but, when the whole universe is declared to be spiritual [i.e. ideal], it is certainly meant to assert that they are far more like us than we think (Moore 1903, p. 433)

According to idealism, as Moore understands it here, the objects of experience are different from how they seem (i.e. independent of experience, distinct from 'us') because they are continuous with our experience of them; like other modifications of subjectivity their *esse* is *percipi* (p. 436). But if, in having an experience of e.g. blue, a subject comes to be aware of

[8] Contemporary discussion of 'sense data' does not always make this clear. For example, Martin contrasts naïve realism's conception of perceptual objects with "subjective entities or qualities" (2002, p. 376) and characterises sense data theorists as claiming that perception consists of "awareness of certain [i.e. these] non-physical or mind-dependent entities" (2002, p. 377). Here, Martin is referring to a later tradition of sense data theorists. Neither claim is essential to the view as expressed by the Cambridge realists – which Martin characterises as a variant on the ideas behind naïve realism. As we've seen, Russell and Moore's argument against idealism, which contemporary naïve realists build on, requires sense data to be constitutively *independent* of the mind's acquaintance with or apprehension of them. Working out just how the earlier tradition understood sense data to be related to ordinary material objects poses some tricky exegetical questions. Over his career, Moore equivocates on the claim that sense data are physically realised in the environment, and Russell often seems to present them as non-physical. But it's at least not obvious that the considerations involved are directly related to their anti-idealist project, in which acquaintance with sense data plays a central role (for more discussion, see Snowdon 2008; Raleigh 2019).

something independent of her – whether this is a property of a sense datum or a distal object, or even a universal – then this vital premise in the argument for idealism turns out to be false. And this result is supposed to follow from the analysis of perception in terms of knowledge and knowledge in terms of acquaintance:

> A sensation is, in reality, a case of 'knowing' or 'being aware of' or 'experiencing' something. […] And this awareness is not merely, as we have hitherto seen it must be, itself something distinct and unique, utterly different from blue: it also has a perfectly distinct and unique relation to blue [...] This relation is just that which we mean in every case by 'knowing'. (p. 449)

Since blueness is not a property of the experience of blue (or, at least, if it were this would not explain its being an experience of blue; p. 450), Moore concludes that it must stand to blue, or to the blueness of some object, in "the simple and unique relation the existence of which alone justifies us in distinguishing knowledge of a thing from the thing known, and indeed in distinguishing mind from matter" (ibid.). In Russell's (1911) terms, it *acquaints* the subject with something independent of her that is blue.

This is what Russell and Moore think they need from perception to secure their case against idealism. And a sense data theory will allow it to play *this* role just as well as naïve realism. The idea that chairs and tables and mountains are real things with properties that correspond to, but are not the same as, the properties of sense data is not a denial of commonsense realism in the way idealism is. So, while Russell and Moore didn't remain naïve realists, they continued to understand perception as putting subjects in touch with mind-independent reality.

This is unconvincing to Travis and Kalderon because they want to push the objection that acquaintance with sense data doesn't put the subject in a better position to avoid a certain kind of sceptical challenge. On their telling, the early analytic sense data theorists, like their predecessors in the empiricist and post-Kantian traditions, lacked the resources needed to answer arguments from illusion, hallucination and conflicting appearances. As a result, their theories "degenerated into phenomenalism" – they fail to measure up to an epistemic criterion for an adequate theory of perceptual experience: that perception is such as to make perceivers knowledgeable about the perceptible properties of the world around them. And these resources needed to meet this criterion would only be provided later at Oxford by Austin (Travis & Kalderon 2013, p. 491).Another way of putting this point is that it is not good enough, from the naïve realist perspective, that perception acquaints perceivers with *just any* mind-independent particulars. The mind-independent reality with which perceivers are acquainted must

encompass the familiar world of material objects and their properties. If unassisted perception leaves us in the dark about these, then realism of the sort naïve realists are interested in has not been vindicated.

The special contribution of the Oxford realists to this epistemic project sets out from "a new conception of philosophical good faith" (Travis & Kalderon 2013, p. 491). This is the proposal that philosophical investigation of the phenomena of ordinary life – like, for example, perceptual experience – ought to be especially attentive to the distinctions made in everyday discourse about those phenomena. Good faith requires that the ordinary language terms used for framing philosophical problems continues to guide philosophers' engagement with them. Cook Wilson expresses this thought as a corrective to philosophers' "tendency to regard the linguistic distinction as the less trustworthy because it is popular and not due to reflective thought" (1926, p. 875, quoted in Travis & Kalderon 2013). The problem with this is the reflective distinctions drawn by philosophers are often idiosyncratic and driven by theory, and so fail to do justice to aspects of the target phenomenon that the philosopher or the theory does not have in view. In contrast, "the experience which has developed the popular distinctions recorded in language is always in contact with the particular facts" (ibid.), which means that it isn't liable to abstract away from perennial and manifest features of the target of investigation.

This methodological proposal recognisably foreshadows a more famous defence of the importance of ordinary language and practice for theorising about perception:

> [O]ur ordinary words are much subtler in their uses, and mark many more distinctions, than philosophers have realised; and the facts of perception, as discovered by, for instance, psychologists but also as noted by common mortals, are much more complicated than has been allowed for [i.e. by philosophers] (Austin 1962, p.3)

This priority given to ordinary language and the everyday distinctions marked by practical common sense expresses an important difference in the approach that the Oxford realists would take to the perception – from Kantian and post-Kantian philosophers but also from the Cambridge sense data theorists. Any analysis of perception as acquaintance with (or apprehension of) sense data is an application reflective thought to the problems of perception, which introduces distinctions that aren't part of everyday talk about seeing or hearing, noticing or recognising, etc. With the Cambridge realists, we've seen that this introduction was set against the background of a sophisticated theory, which dealt with perception as a special case of the mind's basic capacity to apprehend mind-independent reality. While Cook Wilson and

Austin's prescriptions for philosophical good faith don't rule out this approach, they call for careful attention to the way that the established and the new, technical uses of the relevant terms interact.

For the early analytics, sense data are supposed to relate to the perception of distal objects in something like the way observation relates to theory. Sense data provide particular points of evidence for knowledge of facts, which is at a higher level of generality. The Oxford realist response to this – in its fully developed form, as expressed by Austin (Austin & Warnock 1962) – would start out by measuring the use of terms like 'knowledge' and 'evidence' here against the way they occur in ordinary discourse about e.g. knowing things on the basis of seeing them. To do this, as Austin (and contemporary naïve realists) urge is to begin to worry about the decompositional analyses just set out.

Does it make sense to talk of sensation as *evidence* that the subject uses to come to her judgement? According to the Oxford realists, this way of speaking would be a bad faith regimentation of more ordinary talk of perceptual 'evidence', in which perceptions themselves are evidence *for* further claims. One must be in the grip of a theory, they would argue, to cite sensory input as a kind of evidence for one's own perceptual knowledge, for example, that the sun is setting. A more natural response to the question 'how do you know the sun is setting?' is 'because I *can see it*'. Austin's way of making this point is to straightforwardly deny that perceiving something (e.g., watching the sun set, recognising a pig *as a pig*) counts as evidence of that thing being the way it is – although it might count as evidence for some other claim. One has perceptual *evidence* for a judgment, he argues, only when one isn't in the circumstances that would constitute knowledge. To illustrate the point, he discusses the difference between encountering perceptual evidence that some unseen animal is a pig, and being visually confronted by the pig itself:

> If I find a few buckets of pig-food, that's a bit more evidence, and the noises and the smell may provide better evidence still. But if the animal then emerges and stands there plainly in view, there is no longer any question of collecting evidence; its coming into view doesn't provide me with more evidence that it's a pig, I can now just see that it is, the question is settled (Austin, 1962, p. 113).

The question is settled because *seeing the pig,* when the pig is present, puts the subject in the best possible position to know that it is there. While her perceptual beliefs are not directly determined by her perception, since this could be defeated by countervailing considerations (e.g. the subject's belief that she is undergoing a hallucination), the *merely possible* relevance of

these doesn't force the kind of revision of the naïve picture of the ordinary perceptual situation that Russell – and later Moore – thought that it did. When a subject is in propitious circumstances to know that there is a pig in front of her, and perceptually picks up on the pig's presence, then she knows this state of affairs (i.e. that *over there* stands a pig) for a fact. And where there is knowledge of this kind, it makes no sense to speak of evidence.

On this account, the good perceptual case (seeing the pig) takes explanatory priority over a 'neutral' characterisation of the subject's perceptual experience, in which the object might or might not be present. The possibility of hallucination is characterised in terms of failure of the ordinary conditions to hold; but that these conditions *do* typically hold is a condition for the possibility of perception and perceptual knowledge as it is ordinarily understood. It is this ordinary understanding, in which perceptual experience without objective presence is seen as a negative contrast to the 'good' case of veridical perception, that should guide philosophical theorising. The uses that sense data theorists make of terms like 'evidence' and 'knowledge' and the ordinary verbs of perception, are unmoored from and unguided by wider discourse. And if this wider discourse is always in contact with the (relevant) particular facts about perceptual capacities, then proposals driven by intuitions about sense data risk losing sight of these – while, at the same time, illicitly trading on their established explanatory cachet.

c) Contemporary naïve realism and phenomenal character

Contemporary naïve realism draws on both of the realist traditions outlined above. These provide the terms in which the view is standardly presented; and in the work of many naïve realists, the connections run much deeper than this. But there are differences of emphasis that can be used to differentiate available ways of making the naïve realist case. One way to bring these out is to distinguish between the anti-idealist and anti-sceptical motivations for naïve realism, which can be more closely associated with the Cambridge and the Oxford realist traditions respectively.

Generically, contemporary naïve realists defend a claim about the phenomenal character of perceptual experience grounded in a relational account of its metaphysics:

> *Phenomenology:* Perceptual experience is the way it is for the subject because of how it relates her to objects in the environment and their properties.

Naïve realists argue that the way that perception is for the subject accords with some version of the two claims about *transparency* outlined above; these are "the phenomenological echoes of the fact that one is in a state with the functional role experience has" (Martin 2002, p. 391); this

'functional role' depends on experience instancing a *basic* psychological relation between subjects and environmental particulars. At first, this might make it sound like naïve realist is compatible with the claim that perceptual phenomenology depends on *both* representational and relational features of experience – especially if one thought that we should appeal to the former to help explain the latter (cf. McDowell 1994). Naïve realists deny this. They are committed to the view that representation should either play no explanatory role or only a derivative role at the personal level (for discussion, see Philips & French 2023). Explanations of phenomenal character should appeal to features of perceived objects as they in fact are and not as they are represented as being.

Moreover, since naïve realism is a personal-level theory, its characterisations of the functional role of states of perceptual experience do not imply any close connection between these and subpersonal states individuated in terms of their functional interactions with each other. The 'functional role' of an experiential state should be understood in terms of how it makes environmental particulars 'available' to the other states, events and capacities that make up the psychological life of the subject of experience. There are two (non-exclusive) ways in which naïve realists develop this idea:

> *Reference*: Perceptual experience puts subjects in the position to uniquely refer to objects in the environment and their properties.
>
> *Epistemology:* Perceptual experience puts subjects in the position to know about objects in the environment and their properties.

A naïve realism motivated by anti-sceptical considerations will have to focus on *epistemology* but can afford to leave questions about *reference* open. On the other hand, settling these questions in the right way might be sufficient for an anti-idealist version of the project (maybe, e.g., once the capacity for perceptual belief about a mind-independent world is established, further non-perceptual conditions for knowledge come into play). The explanatory priority given to the two claims characterises the two ideal types of naïve realism. Cambridge-style accounts are in the first place concerned to show how perceptual acquaintance makes reference to a mind-independent objects possible; Oxford-style accounts aim to show how acquaintance makes perceivers knowledgeable about the perceptible world.

2) Naïve Realism and Other Contemporary Views

a) Contemporary taxonomies: Martin and Pautz

Naïve realist views of the kind described above, as well as 'mixed theories' that integrate elements of naïve realism into other theoretical frameworks, are an increasingly popular way in which philosophers think about perceptual experience. While, in a fairly recent review, Genone (2016) describes naïve realism as an "underappreciated" view, a textbook published only five years later characterises it as "now one of the dominant theories of perception" (Pautz 2021a, p. 189). Alongside naïve realism, Pautz lists these as sense data theories, internal physical state theories and representationalism. Even if we ignore ways in which the broad frameworks that correspond to these labels might hybridise with each other, it should be clear that they cover a much more variegated and irregular landscape than the neat territorialisation suggested by the list. Framing the debate in this way nonetheless brings into view what is at issue between naïve realists and other contemporary theorists. The distinctions that it tries to capture are between ways of answering the question as to how we should individuate states of the perceiving subject in the most general terms: is it by specifying an array of sense data, a physical state type, a representational content, or a perceived object?

It's worth repeating that the contemporary sense data theories, which are the ones Pautz discusses, are importantly different from the views I've attributed to the Cambridge realists. The objects of perception, according to contemporary sense data theorists, are mind-dependent entities whose intrinsic phenomenal properties account for the experienced character of episodes of perception. We can categorise the older sense data theories together with naïve realism in a matrix derived from yes/no responses to two questions: a) do perceptual experiences depend on the existence of their objects? b) do perceptual objects depend on the occurrence of the experience? (cf. Martin 2006) Filling the other theories listed by Pautz into the same matrix, we get this pattern:

	a) yes	a) no
b) yes	Sense data (contemporary)	
b) no	Naïve realism Sense data (Cambridge)	Representationalism Internal physical state

Table 1.1

This table highlights the shared commitments of perceptual acquaintance views. We can get a rough sense of the role that anti-idealism and anti-scepticism have in motivating this kind of account by thinking about the implication of taking the other possible paths. Answering 'yes' to (a) and (b) invites worries about the slide into perceptual idealism illustrated by Pater (1873) above. The position is marked out by the claim that the objects that a subject perceives do not exist independently of her perception of them. This suggests a symmetrical dependence between perception and object that makes it hard to see how perceptual experience could give rise to thought about a world beyond these immediate objects. But answering 'no' to (a) and (b) seems to prise perception and its objects too radically apart. It entails the possibility of systematic perceptual hallucinations, and so threatens to undermine our grounds for thinking that perception makes subjects knowledgeable about the world around them. Since a 'no' (a), 'yes' (b) response is incoherent, this leaves only naïve realism and related views. These are tailored to reflect the asymmetrical dependence of perceptual experience on the mind-independent world: a fact that they plausibly claim is reflected in its phenomenal character and in the referential and epistemic capacities it affords.

But should we expect a single relation, 'perceptual acquaintance' to underlie the three elements of this attractive picture? Susanna Siegel (2019, p. 349) expresses doubt about the use that contemporary naïve realists make of the legacy of Oxford and Cambridge realism, complaining that these "[h]eirs of acquaintance are recruited to play multiple roles at once, but these roles can be dissociated". In particular, Siegel is concerned to distinguish explanations of perceptual experiences' epistemic properties from explanations of their phenomenal character.

If we understand (a) as a question about the *phenomenal character* of perceptual experiences, table (1.1) accurately categorises the views filled into it. But things look different if we replace it with the following question: a*) does the *epistemic contribution* made by a perceptual experience depend on the existence of its object?

	A*) yes	a*) no
b) yes	Sense data (scepticism)	Sense data (contemporary)
b) no	Naïve realism Sense data (Cambridge) (?) Representationalism	Internal physical state

Table 1.2

Here, categorisation will vary depending on how 'epistemic contribution' and 'dependence' are understood. I've relied on a loosely intuitive notion of the rational (i.e. reason-giving) contribution that perception makes to a subject's beliefs, upstream from other cognitive capacities, and a broad notion of counterfactual dependence. This places contemporary sense data and physical state theories in the right-hand column. On these accounts, perception makes no independent *rational* contribution to the epistemic status of a subject's perceptually grounded beliefs – and so this cannot depend on the existence of a perceptual object.

Representationalist theories, meanwhile, end up on the bottom left – in the same category as the perceptual acquaintance tradition and its heirs. Arguably, they fit here more comfortably than at least Russell's (1912) version of the Cambridge view. According to this, the direct presentation of sense data to the subject is supposed to account for experience's foundational epistemic role, but perceptually grounded knowledge of the environment depends on downstream capacities for inference from these perceptual 'givens'. In contrast, representationalist views explain perception's rational contribution to the epistemic situation of the subject in entirely *perceptual* terms. Attributing content to an agent's perceptions of the world helps explains her successes and failures in action, reasoning, and other capacities. Representationalists are moved by considerations of this kind to individuate perceptual experiences in terms of such content. So conceived, the positive epistemic contribution made by a veridical perception (i.e. a perceptual representation of things in the environment, *as they really are*) is different to, because paradigmatically better than, the negative contribution of misperception. This positive contribution depends, minimally, on the existence of the perceived object.

The arrangement of possible positions given in table (1.2) also brings into focus the kind of scepticism Martin attributes to Hume (Martin 2006, pp. 1, 48-57; see also McDowell 1994, p. 192). If the epistemic contribution of perceptual experience to a subjects' perceptually-grounded beliefs is supposed to depend on the existence of its objects but those objects are not the ordinary material things, events and properties that typical subjects understand themselves to be perceiving, then it is unclear how perceptual experience (i.e., on this view, perceiving sense data) can make perceiving subjects knowledgeable about the real features of their environments. As the same sense data might be *distally* caused by entirely different objects, the grounds we might have for expecting perception to afford this knowledge seem to have been undermined. Avoiding this kind of scepticism while accommodating the attractive claim that

perception makes a rational contribution to the epistemic condition of the subject is a motivation behind representationalist accounts as well as naïve realism.

Filling out a similar matrix for a question about perceptual reference is complicated by the different ways that representationalists, on one hand, and naïve realists, on the other, understand the relation between reference and perception. Naïve realists think that perception is bound up in subject's capacities for reference, and in particular, singular reference to the objects of perception. But they deny that perception intrinsically represents or refers to anything. Representationalists think that perception constitutively depends on the referential capacities that it embodies. One might think of these as proposing 'mixed' theories, that incorporate relational elements into a representationalist framework – or vice versa (for discussion, see Genone 2016). For example, while Burge (2005) claims that externalist representationalist theories of the kind he endorses have been "dominant" throughout the history of philosophy and are "presupposed and relied on in perceptual psychology" (p. 1), he characterises naïve realism and disjunctivism as defective ways of working out the same object-dependence thesis (p. 2; see also Burge 2011). This suggests that, notwithstanding the fundamentally different way they understand the relation between perception and reference, we could categorise this strand of representational theorising together with naïve realism. Both require that a kind of perceptually mediated reference, crucial to understanding the nature of perception, depends on its objects.

A full survey of the logical space probed by this discussion, and by tables (1.1, 1.2), would turn up a number of mixed views. Naïve realism can be characterised (circularly, but informatively), as a kind of pure relationalist view that occupies the bottom left corner of tables (4.1, 4.2) and all similarly constructed matrices, for the right range of questions. While naïve realism about perceptual phenomenology is often the focus of discussion, a question Siegel (2019) raises is how far naïve realist commitments about phenomenology can be maintained without support from intuitions about reference and epistemology that might equally well be accommodated by representationalist accounts.

b) Cognitive science and naïve realism

This is a pressing question because it's often taken as given that cognitive science gives us reasons to prefer representationalist accounts of one kind or another over naïve realism (for discussion, see Epstein 2022; French & Philips 2023). If this widespread impression is correct, and if representational accounts can make sense of the other features of the mind's capacity to be related to things other than itself that acquaintance was originally invoked to explain, then

naïve realism's core claim about phenomenology looks hard to defend. But the case from cognitive science against naïve realism is more difficult to put than many philosophers who feel a tension between them might expect.

As a first pass at articulating this tension, we could highlight two features of mainstream cognitive-scientific models of perceptual processing:

> *Computation*: Perception depends on computational processes that use sensory input to generate representations of environmental objects and their properties.
>
> *Internal realisation*: Perceptual processing is *neuro*computation. It is exclusively realised by processes within the central nervous system of the perceiving organism.

If personal-level perception depends on *computation*, then it depends on representation; if this computation is *internally realised*, then it excludes ordinary distal objects. So, it might be thought, two defining naïve realist claims conflict with assumptions that guide much of contemporary cognitive scientific work on perception.

French and Philips (2023) convincingly make the case that no simple argument from *computation* alone is likely to work.[9] Critics of naïve realism along these lines assume an identity between the computational states involved in perceptual processing and personal-level states of the perceiving subject. But they cannot take this assumption from the cognitive scientific sources they appeal to, since these don't make it.[10] Representational content is posited at the subpersonal level to make sense of the activity of mechanisms thought to underlie personal-level capacities and phenomena; and there is no straightforward argument from this methodological approach to the kind of identity that the simple objection from cognitive science requires. The relational features of perceptual experience that naïve realist accounts of its phenomenal character appeal to might only 'come into view' at the personal level of explanation.

Nonetheless, these arguments require us to provide *some* mapping from the constitutive elements of perceptual experience to subpersonal models of its implementational mechanisms. This allows for the reciprocal adjustment of models at the two levels, explaining the special explanatory relevance of cognitive science for philosophical and everyday thought about

[9] Siegel's (2019) argument is more sophisticated that this outline suggests, and does not rely on the identity assumption French and Philips criticise (see also Siegel 201)

[10] At least, they don't *systematically* make it. The standard formulation in cognitive neuroscience, for example, is that hypothesised representations and observed activity are 'related to' or 'involved in' the relevant personal-level capacity or phenomenon of interest.

perception – and *vice versa*. While straightforward appeals to *computational* subpersonal perceptual processing fail to demonstrate the representational character of personal-level perception, the neurological (and so internal) realisation of these processes might offer greater traction to critics of naïve realism. More sophisticated versions of the challenge from the *internal realisation* of the computational processes underlying perception cannot be ruled out in advance. These leverage the cognitive scientific claim that perception depends on *neuro*computational representations against naïve realist claims about object-dependence. This challenge does not rely on isomorphisms of content across levels. Meeting it involves specifying in general terms how subpersonal neurocomputation bears on personal-level acquaintance.

c) Internal realisation: second pass

The canonical statement of the argument from *internal realisation* in contemporary philosophy of cognitive science is due to Tyler Burge (2005, 2011). The main target of Burge's criticism is 'disjunctivism', and in particular the account put forward by McDowell (1994). But in as far as naïve realists adopt disjunctivism as a strategy for dealing with arguments from hallucination and illusion (Martin 2004; Soteriou 2016), it tells against their proposals too.

The crux of Burge's case against disjunctivism (and so naïve realism) is that it is fundamentally at odds with the science of perceptual psychology, as this has developed in the wake of the cognitive revolution. For Burge, the success of cognitivist models in perceptual psychology (particularly for research on vison), shows that we have "no good reason to doubt that it provides insights not only into the mechanics of perception, but into aspects of its nature" (Burge 2005, p. 9). These successes, however, have been attained by methods that share a central guiding assumption that he labels the *proximity principle*. This states that:

> the kinds of perceptual states that are formed – including conscious state kinds that are the perceivings and misperceivings by individuals – depend purely on (a) the registration of proximal stimulation, (b) the antecedent psychological and physical states of the individuals, and (c) the quasi-deterministic laws of transition between registration of proximal stimulation and the perceptual states that are formed. (Burge 2013, p. 44)

Burges takes it that the truth of the proximity principle implies that naïve realism is false. Imagine a situation in which conditions (a-c) are held constant while the distal objects of perception (i.e. the objective relatum of acquaintance) is changed. Disjunctivists say that this will give rise to different kinds of psychological state; and naïve realists in particular say that it

will give rise to different phenomenal states. But, to the extent that the proximity principle captures the practice of perceptual psychologists, it looks like they would deny this. Perceptual psychology, according to Burge, is committed to the common kind assumption. A presupposition in favour of methodological naturalism is supposed to resolve the issue. Since perceptual psychology is a mature and productive science, philosophers of perception should defer to its guiding assumptions. This means accepting that the perceivings and mispercevings of individuals belong to a shared psychological kind, instances which constitutively depend only on conditions (a-c) of the proximity principle – and this amounts to giving up on both disjunctivism and naïve realism.

At first, this argument might not seem to bring the discussion on further than our first pass at an argument from *computation* against naïve realism. Can't naïve realists simply deny that the psychological kinds that they are interested in are, or directly correspond to, the perceptual states posited by cognitive scientists constrained by the proximity principle? (cf. Campbell 2008, 2011b; French & Philips 2023). Burge stipulates that these include 'perceivings and mispercievings of individuals', but this claim looks open to challenge in a way that doesn't take issue with the proximity principle or any other aspect of cognitive scientific methodology. Such states, as they occur in psychological explanation, are constructs that are posited to explain patterns of behavioural and physiological data that bear on operationalised models of perceptual capacities. Clearly, these have some explanatory connection to personal-level perception; but it is far from clear that they directly correspond to the 'folk kinds' that interest naïve realists (cf. Campbell 2011b). Again, it looks like the critic of naïve realism is relying on assumptions that go beyond those made in the relevant science.

But the assumptions required for this argument to go through are not so easy to dismiss. Take it that Burge is right that at any moment that a subject is having a perceptual experience she, or her brain, or some perceptual subsystem in her brain, is in a neurocomputational state, individuated by specifying values for (a-c). We can accept this while also thinking that *some* psychological states can't be individuated by specifying values for (a-c) or any similarly proximal conditions. The anti-idealist and anti-sceptical arguments canvased above give us reasons to want there to be object-dependent states like this, and the sketch of the contemporary debate over naïve realism that followed on from this discussion suggests that commitments to them are widely distributed across the literature (see also Drayson 2019). In particular, such object-dependent states are often thought to be implicated in subjects' capacities for knowledge of and reference to environmental objects.

I've suggested that the intuitive appeal of naïve realism derives from the thought that the phenomenal character of perceptual experience is bound up with perception's role in supporting one or both of these capacities. Naïve realism offers a simple, attractive way to work this thought out, pitched at the personal level of explanation. The naïve realist proposal is that the subject (and *not* her brain or some subsystem of it) is also in a perceptual state that specifying (a-c) does not uniquely pick out. Individuation conditions for this state also involve the environmental object that causes the proximal stimulation registered in condition (a); and it is these states that fix the phenomenal character of perceptual experience. Even if one can individuate neurocomputational states of the perceiver (or of some perceptual subsystem in her brain) in such a way that they might be put into one-one correspondence with cooccurring temporal parts of her perceptual experience, this mapping does not allow us to explain the phenomenal properties of the latter by appealing to features of the former. We have no reason to expect an *explanatory* correspondence between the two.

Conceding this much, however, leads to problems for naïve realism. Neurocomputational states are characterised by their representational content and (subpersonal) functional role. As Burge conceives these in the perceptual case, their role is to carry content about environmental objects. The idea is that this content makes perceptual states that respect the proximity principle apt to explain the referential and epistemic capacities of subjects. So, the subject might be in object-dependent psychological states while having no specifically perceptual states – phenomenally conscious or not – that constitutively depend on the condition of the distal environment. The possibility of alternative explanations of this kind is what Siegel (2019) is referring to when she claims that cognitive scientific and philosophical development of the notion of representational content allows us to dissociate the functions that Russellian perceptual acquaintance and its contemporary successors unify. And, as I've already suggested, that this should make us hesitant about moving too quickly from claims about the object-involving features of epistemic and referential capacities embodied in or facilitated by perceptual capacities to similar claims about the phenomenal character of perceptual experience.

The challenge to naïve realism that we get from taking Siegel (2019) and Burge's (2005, 2011) arguments together is this: cognitive science explains perceptual capacities by appealing to neurocomputational states that constitutively exclude environmental objects. These have representational content that plausibly realises a basic capacity for singular reference to environmental particulars embodied in perception itself, and they are situated within a functional architecture that enables the subject to be knowledgeable about her environment

partly in virtue of this capacity. That is, neurocomputational representational states in the perceptual system might furnish material for subpersonal explanations of post-perceptual epistemic capacities that do not rest on a questionable identification between them and a perceivers object-involving psychological states, while failing to give environmental objects a constitutive role in shaping perceptual phenomenology. Explanations along these lines offer a way of responding to the sceptical and idealist worries that motivated contemporary naïve realism's historical precursors in a way that makes no appeal to a distinctive personal-level relation of perceptual acquaintance between subjects and environmental objects. A phenomenal role for the acquaintance relation is not ruled out, but it must be motivated. Since neurocomputational states that respect the proximity principle are available for candidate explanations of perceptual phenomenology, the onus is on the naïve realist to carve out a *basic* explanatory role for acquaintance against the backdrop of perceptual neurocomputation.

3) Neurocomputational Naïve Realism

a) Acquaintance and perceptual standpoints

Beck's (2019) neurocomputational naïve realism is a framework for explaining the phenomenal character of episodes of perceptual experience as grounded in a subject's acquaintance with environmental objects, which aims to answer the challenge from *internal realisation* in the terms in which it is put. The prospectus for neurocomputational naïve realism is that it a) allows naïve realists to keep the relational account of the phenomenology of perceptual experience to which they are distinctively committed; b) it respects naïve realists' principle negative commitment against representational accounts of this relation; c) it allows them to take advantage of whatever theoretical benefits accrue from (a) and (b) while, d) accommodating the pull towards neurocomputational explanations of perceptual phenomenology. The basic proposal is that relational phenomenological states of the subject don't compete with neurocomputational states of her perceptual system; they incorporate them.

Beck's argument has three main stages. First, perceptual acquaintance is analysed as a three-place relation, between a subject, S, an object, O, and a way, W, in which that object appears to the subject (cf. Campbell 2002). Beck notes that analyses of this kind have a distinctive advantage: they account for the 'external-directedness' of perceptual experience. According to Beck, this external-directedness is what claims that perception is $transparent_{Ext}$ to environmental objects are supposed to capture. Like all naïve realist models, such three-place approaches also face a challenge. They must account for perception's 'internal-dependence'. This is what Beck thinks we should take from the argument from *internal realisation* outlined

above. While external-directedness is built into any naïve realist analysis of perception, a distinct advantage of the three-place model is that it offers a way of clearly articulating what it would take to simultaneously accommodate internal-dependence.

Second, Beck distinguishes between 'selectionist' and 'non-selectionist' versions of naïve realism. Selectionist naïve realists are committed to an implausibly strong claim about *transparency*$_{Ext}$ and so are vulnerable to arguments from *internal dependence.* The strong claim is that the phenomenal character of an experience is exclusively determined by properties of the perceived object: the way that an object appears to the subject amounts to the *selection* by perceptual processes of a set of its mind-independent perceptible properties. These processes can make no positive contribution of their own to the phenomenal character of experience. Beck argues that selectionist naïve realism puts relational and neurocomputational explanations of perceptual phenomenology in conflict, and he argues on empirical and methodologically naturalistic grounds that we should – in case of such conflict – prefer neurocomputational explanations.

Third, Beck outlines his positive alternative account: neurocomputational states of the perceptual system determine the way objects perceptually appear to the subjects of perceptual experience. Since this is a non-selectionist proposal, the 'way' objects appear involves a positive contribution from neurocomputation to overall phenomenology. They determine the 'appearance properties' of the experience. Nonetheless, while appearances are grounded in neurocomputational representations, they do not exclusively determine *perceptual* phenomenology. This further depends on the objects of perception *appearing to* the subject – that is, on the O *appears* W *to* S relation being instanced.

Perceptual phenomenology, so conceived, constitutively depends on a non-representational relation between subjects and objects that is distinct from, but incorporates, appearance properties fixed by internally realised neurocomputational states. This is how neurocomputational naïve realism fulfils items (a), (b) and (d) of its prospectus. But I'll suggest below that it leaves open a worry about (c). If all neurocomputational naïve realism delivers is a stipulated difference in the metaphysical grounding conditions of perception and non-object involving states of sensory experience, Beck's positive proposal will be unmotivated. Absent another way of working out non-selectionist naïve realism, it would look like we should default to purely neurocomputational – and so representationalist – accounts of the phenomenal character of perceptual experience.

While Beck (2019) only explicitly defends a proposal that meets 'minimal criteria' for naïve realism (p. 608), the neurocomputational naïve realist framework that it illustrates is supposed to be hospitable to a wide range of existing versions of the view. In particular, he describes the idea behind the framework as already present in Martin (1998), while his treatment of perceptual acquaintance as a three-place relation is a simple extension of the model proposed by Campbell (2002).

Beck is attracted to the three-place relation model of perceptual acquaintance for a similar set of reasons to Campbell's. Introducing a third, multi-component factor – call this a perceptual 'standpoint' (cf. Campbell 2002) or 'viewpoint' (cf. Martin 1998) – allows naïve realists to make sense of perceptual variation and constancies in a plausible, simple way. The manner in which a perceiver is related to an object (e.g. her physical position vis-à-vis that object, the perceptually relevant conditions under which it is encountered, her 'mode of engagement' with it) intuitively all make a difference to perceptual phenomenology. The three-place relation model allows naïve realists to bring these together under a single heading.

Not all naïve realists frame the account in this way, but it's clear from the way in which Beck introduces the model that it isn't intended to tie neurocomputational naïve realism to the full set of commitments and motivations that I've associated with Cambridge-style naïve realism above. Its role in the framework is limited to the way it can be used to capture the "important idea" (Beck 2019, p. 624) expressed in this passage from Martin (1998):

> [T]o have an experience is to have a viewpoint on something: experiences intrinsically possess some subject-matter which is presented to that viewpoint. To understand such experience and what it is like, one has to understand the viewpoint on that subject-matter, and hence also to attend to the subject matter *as presented to the viewpoint*. (pp. 173-174; quoted in Beck 2019, emphasis added)

For all naïve realists, differences in the objects of perceptual experience are sufficient for differences in the phenomenal character of those experiences. Beck identifies this as the "minimal disjunctivism"[11] entailed by naïve realism. This is a weak version of *transparency*$_{Ext}$, and so securing this entailment is plausibly a minimal criterion for naïve realism (ibid.). But what's particularly important about this passage for Beck's framework is that it makes space for the following question: "why cannot the *ways* in which things are presented in experience make a difference to what experience is like, in addition to what is presented?" (Martin 1998, p. 175).

[11] Cf Burge 2005 discussion of disjunctivism

Beck reads this as a rhetorical question about the basis for perception's *transparency*$_{Ext}$ to its environmental objects. Accordingly, what's being suggested is that both the objects presented and their modes of presentation are components of the overall phenomenal character of the subject's perceptual experience. So, while a difference in presented object entails a difference in overall phenomenal character, two experiences of the same object can be phenomenologically different just because the way in which that object is presented differs. In particular, Beck's way of reading Martin (1998) is consistent with thinking that the way in which an object is presented to a subject makes its own positive contribution to the phenomenal character of a perceptual experience. It is not just an additional explanatory factor that analyses of perceptual experience should take into account, it makes what we might call an 'additional' difference to phenomenology. This picture of how the components of perceptual phenomenology interact is crucial for the distinction Beck wants to make between his proposal and selectionist versions of naïve realism.

b) Selectionist naïve realism

According to Beck (2019, pp. 609-610), selectionist naïve realism is marked out by two claims:

> [P]erception has the phenomenology it does completely in virtue of the items perceived in the perception.
>
> [S]tandpoint conditions serve to select which of the many items in the subject's environment the subject perceives.

While the first claim might seem to cut against the motivations for introducing a third relatum of perceptual acquaintance, there are still reasons selectionist naïve realists might want to frame their account in terms of the three-place relation model. Because of selectionism's distinctive commitments, however, these get expressed somewhat differently – and this corresponds to a significantly different understanding of the role and nature of the perceptual standpoint conditions that are supposed to underlie it. We can get clearer on how the version of naïve realism Beck finds in Martin (1998) differs from selectionism by contrasting how this third relatum is conceived on the two accounts.

As Beck points out, the relation o *appears* W *to* S is intimately connected to the relation S *perceives* o. In particular, he highlights that o's *appearance* to S entails S's *perception* of o. This is what secures the connection between phenomenal character of experience (fixed by the first relation) and the objectivity of the perceptual relation. So, the motivation for using a three-place relation to model perceptual phenomenology comes from taking perceptual objects to include

material things that might appear in different ways to the subject. Since the ways o can appear aren't fixed by intrinsic features of the subject or the object, the third relatum is introduced to cover this extra variability. The idea is that even if the phenomenology of perception is entirely determined by its objects, we should want to take account of the way in which they do this.

If o *appears* W *to* S, it also follows that S *is appeared to in* W and that o *appears* W. Can these two-place relations be instanced separately from each other? Selectionism says that they can't. According to selectionism, the ways in which a subject can be appeared to are tied up with the properties of the object that the subject perceives. Standpoint conditions select which perceptible properties (e.g., elliptical-from-here, red-in-these-conditions) shape the subject's experience, but they can't make any independent difference to it. This looks intuitive, if we think of standpoint conditions as determined by how the perceiver is physically oriented in her environment. Physical conditions like this determine perceptual phenomenology in all sorts of ways, but it's clear that the contribution they make can't be detached from the concrete perceptual relations they are involved in since these constitutively involve environmental objects. According to selectionism, all standpoint conditions should be understood in this way.

According to non-selectionist naïve realism, there are cases in which instances of S *is appeared to in* W and o *appears* W can be pulled apart. Perceptual appearances are sometimes detached from their objects. 'Detachment' in this sense is can be introduced via considerations about the adverbial modification of perception. Under certain physiological conditions (e.g., tiredness, drunkenness) a subject might see *blurrily* (French 2014) or under other, psychological, conditions she might listen *attentively*. Whatever phenomenological difference is made by these conditions, it is uncontroversial that they don't involve an objective change in the environmental objects of perception. It looks plausible, then, that they should be accounted for as a positive phenomenal contribution made by the relevant condition. If standpoint conditions are conceived on the model of these (and this interpretation is assumed), then it looks like we can make sense of a way in which a subject is appeared to by a perceptual object, without attributing to that object an intersubjectively available appearance property.

Beck (and Martin, as Beck reads him) radicalises this proposal. For him, appearance properties are always properties of the subject:

> Appearance properties […] correspond to Martin's [op. cit.] ways of presentation. To have an appearance property is to be appeared to in a certain way. For example, you can be appeared to in a roundish way or a blueish way. (Beck 2019, p. 629)

In a footnote, Beck clarifies that this use of 'appearance properties' is to be distinguished from another use, current in the naïve realist literature, according to which "appearance properties are not properties subjects can have, but […] are ways that objects and scenes appear *simpliciter* (as opposed to the way they appear to a specific subject)" (ft. 24; see, e.g. Kalderon 2011; Genone 2014). Appearance properties in Beck's sense are fixed by instances of the S *is appeared to in* W relation and are independent of instances of o *appears* W in the more established sense. This is a crucial stage in Beck's argument, since he thinks that only this strong version of non-selectionism can accommodate the *internal realisation* of perceptual processing.

c) Neurocomputation and appearance properties

The particular impetus for Beck's rejection of selectionism and his proposal of a neurocomputational alternative comes from a series of articles by Pautz (2013, 2014, 2023). If we take it that Burge and Siegel's arguments establish a candidate class of subpersonal explanations of perceptual phenomenology that do not give any role to perceptual acquaintance, Pautz can be seen as making an empirically grounded case that we ought to prefer explanations of this kind. Once we see that neurocomputational states are available for subpersonal explanation of perceptual experience, the lesson that Beck thinks we should draw from Pautz's survey of the relevant empirical literature is that naïve realists cannot remain committed to selectionism.

Pautz's argument rests on a pattern of empirical results that he interprets as setting up correlations between the qualitative properties of experience (as captured by behavioural measures) and internal neurophysiological processes on the one hand and external objective properties of the stimuli on the other. The challenge posed by these to naïve realists is independent of conditions that involve illusion or approximate the philosophical conception of hallucination, since

> [p]sychophysics has shown that, even in normal cases, qualitative similarity is very poorly correlated with external physical similarity. In short, the typical situation is that there is "good internal correlation" even while there is "bad external correlation". (2017)

The basic argument is as follows. Naïve realists claim that perceptual phenomenology is shaped by the objective properties of the environmental objects that the subject is perceptually acquainted with. But phenomenology can vary without any change in these – and these can change without any variation in phenomenology (at least, in as far as phenomenological changes

are reflected in behavioural measures). Pautz takes this to show is that "looking for the basis of qualitative character in the world is looking in the wrong place; we have to be looking rather at the brain." (ibid.) To the extent naïve realists deny this, they "neglect the scientific facts" (ibid.) as revealed by this pattern of results.

The results are not only that, under a range of conditions, perceptual scientists report a double dissociation between phenomenology and objective properties of perceived objects. The processes that do correlate with contrasts in phenomenological character are identified in ways that appeal to their computational – and so representational – functions. This allows for functional decomposition of neurophysiological processes, and so can be used to identify patterns in the physiological data that underly the strongest correlations with reported phenomenology. These processes are attributed content that makes sense of their functional role – i.e., representing features of perceived objects. So, explanations framed in this way can be used to account for both the phenomenal external-directedness and the internal-dependence of perception, in a way that screens off its environmental objects.

While Beck identifies a few strategies selectionists might use to respond to this argument, each involves at least one of two general problems. They are either inadequate to the empirical evidence for *internal realisation* or they rely on a wider framework of claims about perceptual acquaintance that the availability of neurocomputational explanation puts into question. The distinctive explanatory role that selectionist naïve realism gives to acquaintance cannot be maintained independently of these; but in the context of alternative explanations furnished by cognitive science, commitment to them seems to require prior commitment to naïve realism.

The good news for naïve realists, however, is that they don't have to be selectionists; and non-selectionist naïve realists can adopt a simple strategy to accommodate the results Pautz appeals to. The first step of this strategy is concessive: naïve realists should accept that these results "suggest that subjects have the appearance properties they do completely in virtue of their neurocomputational properties" (p. 610). This means they should identify the standpoint conditions of the perceiving subject with the neurocomputational processes that underlie her perceptual experience. Consequently, these turn out to be fully *internally realised.*

This does not tell against naïve realism *per se*, because:

> to say your neurocomputational properties determine your appearance properties is *not* to say your neurocomputational properties completely determine your perceptual phenomenology. (pp. 623-624)

Non-selectionist naïve realists can accept that neurocomputation determines appearance properties because they can see these as separable from environmental objects, just as the perceiver's neurophysiological properties are. They can then use the positive contributions these make to the overall phenomenal character of perceptual experience (by realising the subject's neurocomputational standpoint conditions) to account for the pattern of empirical results correlating aspects of perceptual phenomenology with neurocomputation. Nonetheless, the overall character of a perceptual experience depends on the O *appears* W *to* S relation.

> Appearance properties are among the relata of this relation, and distinct from it […] [A] ball's perceptually appearing round is one thing, and your being perceptually appeared to in a roundish way is another. Your neurocomputational properties determine the latter. The former requires you to perceive the ball. (p. 625)

So, a difference in the object perceived will always be sufficient for a difference in perceptual phenomenology. This is because it will involve a change in one of the relata, O, on which the overall phenomenal character of a perceptual experience depends. For the same reasons, changes in standpoint conditions are also sufficient. And since these are realised by variations in neurophysiological properties of the subject (or between different subjects), they are detachable from the objects of perception in a way that allows non-selectionist naïve realists to make sense of the empirical data that Pautz appeals to.

d) Neurocomputation and transparency

Two worries, internal to naïve realism, about this proposal concern how far it really fulfils minimal relational and anti-representational criteria. How far does it satisfy items (a) and (b) of the prospectus on which Beck advances it? Under heading (a), a critic pressing this kind of worry might note that Beck's characterisation of Martin (1998) seems to be in tension with the latter's influential, later statements of naïve realism (e.g. Martin 2002, 2006). The most salient point of tension centres on Beck's characterisation of the 'positive contribution' made by appearance properties to perceptual phenomenology. Beck notes a divergence between his and Martin's (2004) account of the phenomenal character of hallucinations, since he thinks that the positive phenomenal properties of shared appearances can be used to account for phenomenal similarities between perceptions and hallucinations. For Martin, positive characterisation of hallucinatory appearances of this kind are a unifying assumption behind the key arguments *against* naïve realism. While Beck acknowledges that the neurocomputational framework "must eventually face" (p. 627, ft. 27) this point, if Martin is right then Beck's proposal radically undercuts a main line of argument for naïve realism.

These tensions also have a bearing on the interpretation of Martin (1998) that Beck pegs his account to. The question quoted by Beck comes as Martin is attempting to "[lay] out a common framework for the debate about the nature of experience and perceptual appearances" (ibid.). The strong transparency claim he is considering is Price's (1932) 'insistence' on "the diaphanous nature of experience: namely that sameness and difference of phenomenal properties just are sameness and difference in presented elements." (Martin 1998, p. 175). We're called back, in this passage, to the Cambridge realists' discussion of *transparency*$_{Ins}$. Moreover, it follows directly from a stipulation on 'phenomenal properties' and 'phenomenal character' that we "use these terms strictly to apply only to experiences and their properties and not to the objects of experience and the properties they appear to possess." (p. 174). There's therefore a case for reading 'things presented in experience' as things presented *by* experience – that is, what experience appears to present subjects with from the first-person perspective – independently of a commitment to the existence of perceptual objects of any particular kind (or of any at all). This would suggest that the 'common framework' that Martin sets out is supposed to be shared between 'intentionalist' (i.e. representationalist), sense data, and naïve realist accounts. These all offer different explanations for *transparency*$_{Ins}$, but Beck takes it as read that we should move from this to the approach distinctive of naïve realists: explaining introspective transparency of experience by appeal to the *transparency*$_{Ext}$ of perception to environmental object.

This reading is what Beck uses to justify assimilating Martin's 'ways in which things are presented' to 'W' in the O *appears* W *to* S relation. This relation is supposed to be what underlies the transparency$_{Ext}$ of perception given the inadequacy of 'realist' sense data views like Price's, but it's not obvious that these concerns about the objective metaphysical structure realised in perception are what Martin has in focus in the passages that Beck is drawing on. Following on from them, the target that Martin most immediately engages with is an 'adverbial' account, on which perceptual experience is "merely […] a mode of being affected by the world" (p. 178)[12] in a certain way (e.g. roundishly or blueishly). These fail to capture the *phenomenal* external-directedness of perceptual experience. They deny that perceptual experience is *transparent*$_{Ins}$ in the way Moore and Price thought it was because they reject the explanation these gave in terms of its *transparency*$_{Ext}$ to objective sense data. Martin sees this as confused – adverbial theorists fail to notice that we can give other accounts of *transparency*$_{Ins}$ that are independent

[12] Earlier Martin summarises the "key idea" shared by adverbial approaches in more detail as the proposal that "we should principally think of our experiences as effects upon us by the environment; effects which have a distinctive qualitative character, and which are such that they bring about beliefs about the environment." (p. 165). This characterisation captures some of the epistemic consequences of the view, which are significantly what is at issue in his rejection of it.

of claims about *transparency*$_{Ext}$ – and ultimately "unintelligible" (p. 168) as an account of perceiving subjects' introspective knowledge about the nature of their own experiences.

This reading of Martin (1998) raises worries about how far claims about the *transparency*$_{Ins}$ of perceptual experience introduced in the context of an argument against adverbial conceptions of its phenomenal character can be transferred over to Beck's case against selectionism. But it also generates difficulties for Beck's positive proposal, since the way that this characterises appearance properties is strikingly similar to the advebialist proposal about overall phenomenology. Compare for example C. J. Ducasse:

> 'blue', 'bitter', 'sweet', etc., are [not] names of objects of experience nor of species of objects of experience but of species of experience itself. What this means is perhaps made clearest by saying that to sense blue is then to sense bluely, just as to dance the waltz is to dance 'waltzily' (i.e., the manner called 'to waltz') to jump a leap is to jump 'leapily' (i.e., in the manner called to leap) etc. (Ducasse 1942, quoted in Martin 1998)

And Beck:

> When the x *perceptually appears* W *to* S relation is instantiated in the ball's perceptually appearing round to you, the constituents of this relation are yourself, your property of being appeared to in a roundish way, and the ball (Beck 2019, p. 625).

The key difference between the two views is that while Ducasse suggests that we use instances of S *is appeared to in* W to type perceptual experiences, Beck thinks we should resist this conclusion. For Beck, perceptions are more than appearances, but he seems to agree with Ducasse that appearances themselves are simply adverbial modifications of subjectivity. This is just what thinking of them as detached from their objects in the way he proposes amounts to.

This view can't be ascribed to Martin. The central argument of Martin (2002), for example, is that reflection on the *transparency*$_{Ins}$ of appearances shared between veridical perception and perceptual imagination gives us reasons to prefer naïve realism over intentionalist accounts that might otherwise seem to accommodate the 'first-order' transparency of ordinary perceptual cases equally well. It's difficult to see how Beck could reformulate this argument within his framework, since this ties 'external-directedness' to full instantiations of the full o *appears* W *to* S relation. So, the consequences of the difference between Beck's and Martin's accounts would extend beyond the way they deal with hallucinatory experiences. On Martin's (2002, 2006)

general account, perceptual appearances as such are apt to make subjects knowledgeable about perceptual objects, and this is reflected in the fact that introspection on perception is sufficient to reveal features of the mind-independent world. According to Beck, however, introspectable appearances are fixed by instances of S *is appeared to in* W. This blocks off Martin's way of making the epistemological case for naïve realism.

Of course, this isn't the only way to make this case. Instances of o appears W to S guarantee that S perceives o, so Beck might argue along the lines I've drawn out in Campbell (2002) that perceptual acquaintance fixes reference to mind-independent objects. In the right conditions this can make subjects knowledgeable about them, and the standards conditions that perceivers find themselves in are conditions of this kind.

Here, however, worries about item (b) come into consideration. The simplest interpretation of neurocomputational states appealed to by subpersonal explanations of perceptual experience is that their representational content specifies the objects of perception. This seems to be assumed by Pautz, where he argues that the pattern of results he discusses supports representationalist account of perception (rather than e.g. physical state identity). So, for Pautz, it is these contents that shape perceptual consciousness, and not the environmental objects which in good cases causally give rise to them. If Beck thinks of neurocomputations underlying appearance properties as involving content of this kind, it looks like naïve realism's distinctive commitment to the *basicness* of perceptual acquaintance is threatened. In general, one could appeal to the contents of appearances to explain perception's role in the psychological life of the subject; these would, for example, offer a non-object involving way of accounting for the epistemic contribution of perception.

Beck doesn't make this commitment. But he has very little to say about the way in which a subject's neurocomputational properties make their distinctive positive contribution to overall perceptual phenomenology. We should assume, however, that they are not supposed to do this in a way that would this undermine the explanatory basicness of the O *appears* W *to* S relation – at least for overall phenomenology, but in light of the above discussion, also (ideally) for the epistemological situation of the perceiving subject.

A suggestion for how perceptual neurocomputation might facilitate conscious access to and knowledge of the external world without representing it as such comes from developing a suggestion already tabled by Campbell (2002). Think about the metaphor Campbell uses to motivate a conception of that perceptual experience as $Transparent_{Ext}$. In, an unstable medium requires constant delicate adjustment to remain transparent, with these being guided by a

correspondingly fine sensitivity to the environment. One might think of the subpersonal representations involved in perceptual processing as having a similar role – guiding state transitions in the neurocomputational apparatus of perception so as to sustain its transparency to environmental objects. On this model, naïve realists could accept that perceptual processing depends on internally realised computation over subpersonal representational content, while rejecting a representationalist account of the corresponding personal-level capacity (French & Philips 2023).

One might expect Beck to endorse a proposal of this kind. It allows subpersonal neurocomputational process to play a robust explanatory role in the phenomenology of perceptual appearances since it suggests that these depend entirely on perceptual processes individuated in terms of their functional role. It is plausibly adequate to the empirical evidence discussed by Pautz because just those patterns of neurophysiological data that can be picked out by assigning contents that reflect environmental objects and properties can be identified on the assumption that the relevant areas and networks facilitate access to those same objects. Moreover, as we'll see in the next chapter, approaches along these line are current in cognitive science: many enactivists reject the assumptions that Pautz takes as given.

But there are strong reasons to doubt that Beck can adopt this strategy. First, on many plausible ways of developing it, it will involve taking standpoint conditions to include more that neurocomputational processes. Since subpersonal perceptual representations conceived of on this model radically underdetermine their objects, researchers in this tradition appeal to sensorimotor and other environment involving capacities to explain how these processes make environmental objects perceptually available. To the extent that these involve environmental feedback, they depend on 'extra-dermal' processes that Beck excludes by stipulation from his account of the perceptual standpoint. It is central to the neurocomputational naïve realist framework that "your neuro-computational properties *completely* determine the ways you are appeared to" (p. 625, my emphasis). This is not to deny that non-neurophysiological processes are irrelevant to the standpoint the subject occupies, but it entails that they at most "play a causal role in the instantiation of appearance properties" (Beck 2019, ft. 26). Constitutively, neurocomputation alone fixes perceptual appearances.

Even if a proposal along these lines is constrained to exclude non-neurophysiological or extra-dermal elements, however, it might still fail to be a non-selectionist account according to Beck's criteria. This is because of the way he understands the positive contribution of the standpoint to perceptual phenomenology. This is in tension with perceptual *transparency*$_{Ext}$ as it might be

conceived on the model just sketched, because it still looks like what the neurocomputational components of the standpoint only make available a world of mind-independent perceptual givens. Beck worries that such accounts run afoul of Pautz's argument, since they can offer no plausible account of what produces a phenomenal contrast between two neurocomputationally different subjects who perceive an (quantitively) identical object under (qualitatively) identical conditions. At this point, however, we should interrogate the basis for the argument. This last case is not introduced directly from the empirical literature. Instead, it enters the discussion via a thought experiment derived from these results. So, this part of his argument from internal-dependence looks open to question in a way the rest of it isn't. It is not clear how it gains support from a presumed commitment to methodological naturalism. But if the argument goes through nonetheless, it is also unclear what grounds Beck has for the explanatory *basicness* of perceptual acquaintance to perceiving subjects' capacities to refer to and know environmental objects.

e) Minimal neurocomputational naïve realism?

While this all suggests that neurocomputational naïve realism will be more difficult for standard naïve realists to accept than Beck presents it as being, it doesn't yet give us independent reasons to reject it. We might think that standard naïve realism is ruled out by Pautz and Beck's arguments against selectionism, but still be attracted to a relational metaphysics of perceptual phenomenology. The framework Beck proposes might be able to deliver on this, even if the ways that it meets items (a) and (b) end up cutting out much more of the naïve realist programme than it proports to.

The arguments made in this section don't rule this out. But they give us reasons to doubt that the framework's way of combining naïve realist claims about the basic metaphysical structure of perceptual experience with mainstream cognitive scientific models of its neurocomputational implementation is well motivated. The explanatory framework to which neurocomputational explanations belong constitutively excludes environmental objects from perceptual processing. Neurocomputation is essentially internally realised. This means that introducing the neurocomputational properties of the subject as potential relata of perceptual acquaintance severely limits what elements of the subject's psychology can be explained in terms of a basic relation of perceptual acquaintance – it looks like there are always neurocomputational appearance properties ready to hand to do the same job without introducing objects of acquaintance. But, conversely, the explanatory framework in which distinctive elements of naïve realist explanation are introduced presuppose a rejection of appearances of the kind that Beck introduces. So, it is questionable if they can be used to characterise the standpoint conditions

of the perceiver, while letting perceptual standpoints play their original explanatory role in the theory.

4) Conclusion

Naïve realism offers an attractive picture of how basic capacities of perceiving subjects are grounded in these subjects' relations to environmental objects. It captures anti-sceptical and anti-idealist intuitions that we'll see below are also at work in Hurley's arguments about the perspectival unity of perceptual experience. But it also involves commitments that are at odds with or orthogonal to the project I'll map across the next two chapters. To the extent that naïve realism is a radically anti-representationalist proposal, and to the extent it rejects the action-oriented framing that Hurley's discussion relies on, the two accounts are at odds. By outlining the central elements of the naïve realist picture, and identifying the motivations involved in them, I've tried to lay the groundwork for a later assessment of how far the sensorimotor acquaintance view can share in the attractive naïve picture of the personal level (§5).

When relational accounts of personal-level perception are mapped onto neurocomputation, an assumption that elements specified at each level should somehow correspond to each other can reoccur. We might think that internal and external components at the personal level should map to neurological and non-neurological components of a subpersonal-level explanatory model. I have suggested that something like this assumption is in the background of Beck's (2019) development of a 'neurocomputational naïve realism'. Beck argues that in order to meet the challenge from internal realisation, naïve realists should accept that perceptual appearances are determined by subjects' occurrent neurocomputational processes, while the overall phenomenal character of veridical perception is determined by the acquaintance relation. Perceptual appearance, like neurocomputation, is internally realised.

This leads to trouble, since this move requires a positive characterisation of the phenomenal content of appearances in terms of the representational content of the corresponding neurocomputation. Moreover, if phenomenal states carrying environment-disclosing content arise non-relationally, it becomes difficult to see why capacities for perceptual reference depend on the acquaintance relation. This leaves distinctive naïve realist claims linking overall phenomenal character to acquaintance dependent on the epistemological consideration that only acquaintance with an object can make the subject knowledgeable about its perceptible properties. But the resources for an Oxford-style rejection of phenomenalism are cut off by the foregoing argument. So, it is not clear why subjects of perceptual experience, as conceived by

neurocomputational naïve realism, must be knowledgeable about more than the subjective properties of their experiences.

In this chapter, I've claimed that an objection along these lines succeeds. But this negative argument builds towards my positive case for the sensorimotor acquaintance view. Understanding the motivation for Beck's account has helped show why naïve realists, and other philosophers of perception, should not rely on a generic appeal to the personal/subpersonal distinction to respond to arguments from cognitive science. Certain ways of making this distinction entail an implausibly strong claims about the autonomy of the personal level, others recapitulate the problematic assumptions behind phenomenal internalism. Reflecting on what goes wrong in Beck's account, I will argue, shows that *internal dependence* should be modified. Neurocomputation is central to, but not coextensive with perceptual processing. Sensorimotor processes guided by internally realised representations play a fundamental role in integrating the internal and external relata of perceptual acquaintance. For this account to be available, we have to specify and reject the problematic assumptions it embodies. Showing how this might be done is the goal of the next chapter.

2

Hurley, Dennett and the Philosophers of the Personal Level: Psychological Interpretation and Subpersonal Explanation

0) Introduction

The previous chapter left us with the problem of reconciling two levels of explanation. At the personal level, we have compelling reasons to adopt a relational account of perceptual experience. Trying to make sense of how perception affords subjects of experience a basic mode of psychological access to a shared objective world and to characterise the distinctive phenomenal character associated with this mode of access, it looks like we make progress when we understand it as a kind of acquaintance with environmental particulars – with perceived *objects*. The objects of experience and the character of experiences of those objects seem to be systematically related. And since the environment is composed of objects that can be the shared focus for the experiences of more than one person, a relational, object-oriented approach is well-placed to account for both the objectivity and subjectivity of experience: multiple subjects can be differently related to the same object, their experiences constituted by distinct perspectives on the things that make up their shared environment. Moreover, independently plausible claims about the phenomenal particularity and singular content of object perception fit naturally with a theory that makes sense of perceptual experiences as constituted by relations to particulars.

Naïve realists aim to capitalise on these positive features of the object-relational approach. But a persistent criticism has been that they fail, as Burge has put it, to make their proposals compatible with the 'serious science' of perceptual psychology (Burge 2005, 2011, 2022, pp. 401-405; see also Pautz 2014, 2023). As some respondents to this charge note (see Epstein 2022; French & Philips 2023), 'compatibility' might not in fact be what's at issue. Many of the explanatory aims that motivate naïve realism, as well as the solutions naïve realists offer, can be given a naturalistic treatment tailoring them for compatibility with mainstream cognitive scientific models of the mechanisms underlying perceptual experience. Meanwhile, philosophical 'internalist' theories about the personal-level states they support (as well as their epistemic and phenomenal features) is something that, plausibly, the current empirical framework is silent on. But while some defences of naïve realism rest on the *compatibility* of naïve realism and neurocomputational explanations of perception, Burge's challenge can also be understood as drawing attention to patterns of explanatory dependence in the relevant

scientific work that sharpen and reemphasise traditional objections to the position. As Siegel (2019) develops this response, the challenge naïve realists really face is to justify their characteristically ambitious package of 'realist' claims – about objects and their properties, about the epistemic and semantic import of experience, and the tie between these features and its phenomenal character – over alternatives that bring a lighter load of additional commitments to bear on the cognitive scientific picture. In summary: it looks like we're faced with an explanatory disconnect between the neurocomputational mechanisms that realise perceptual experience and the nature of the psychological phenomenon, as conceived of by naïve realists.

We've now seen that one recent attempt to fit naïve realism into this picture does not justify itself in light of the costs emphasised by Burge, Pautz and Siegel. Although designed to respond to their objections, Beck's (2019) 'neurocomputational naïve realism' fails to motivate the additional commitments it brings with it, and so opens itself to the modified version of the anti-naïve case just outlined. I think that part of the reason for this is that in formulating his proposal, Beck accepts not only cognitive science's results, methods and models, but also a philosophically contentious *isomorphism assumption* about the relation between subpersonal models and psychological interpretation at the personal level. This is the problematic assumption that the previous chapter alluded to, and part of the job of this chapter is to spell out what making it commits one to.

Isomorphic interlevel explanations depend on one-one correspondences between elements of a psychological phenomenon and components of mechanisms underlying it. They form the natural backdrop for the internalist picture challenged by naïve realists, since they exclude mechanism-external relations and objects. But if we have independent reasons to accept a relational account of the personal level phenomenon, shouldn't this motivate a different account of its interlevel explanation? This question motivates Hurley's (1998, 2008) proposal of a 'two-level interdependence' view. And recent attempts to integrate her and other sensorimotor enactivists' models of the subpersonal basis of perception with naïve realism's personal-level acquaintance relation rely on a similar vision of the relation between subpersonal and personal levels of explanation (see, e.g., Ward 2016, 2023; Kirchoff & Kiverstein 2019; Raleigh 2021). Evaluating how far sensorimotor and naïve realist models of acquaintance can be reconciled – with each other and with the mainstream of contemporary cognitive science of perception – will therefore depend on how we answer this question. Doing this, meanwhile, requires us to get clear on what the distinction between personal and subpersonal levels amounts to.

This chapter traces the notion of subpersonal explanation back to its roots in Dennett's foundational contributions to the philosophy of cognitive science (Dennett 1968, 1969, 1978), adopting an approach pioneered in Hurley (1998) and Bill Child (1994) and Jennifer Hornsby's (2000) readings of this work. These emphasise the continuities in Dennett's proposal of a subpersonal mode of explanation with the 'ordinary language' movement, noting especially the influence of Ryle, Anscombe and Wittgenstein on Dennett's original formulation of the personal/subpersonal distinction.[13] I follow Hurley and many contemporary philosophers of cognitive science (see, e.g., essays collected in Muñoz-Suárez & De Brigard 2015), in placing Dennett's proposal at the centre of a widely adopted model of cognitive scientific explanation. I follow Hornsby and Child in associating this with an 'interpretationist' methodology in the philosophy of mind, grounded in the work of the ordinary language philosophers. Subpersonal explanation, I argue, proceeds by content-ascriptive mechanistic decomposition of psychological phenomena. Content ascription is guided (at both levels) by interpretive norms of holism and environmental appropriateness. Appealing to aspects of the environment and prior mental phenomena, these provide 'horizontal' explanations that ascribe content to states of target of psychological interpretation – I'll call these 'interpretability criteria.' Appeal to contents at the level of mechanisms are sensitive to a narrower range of contextual conditions but can themselves feed back into personal-level interpretation. That is, they provide 'vertical' interpretability criteria for the whole agent. Since interpretationists think that criteria of interpretability are constitutive of psychological phenomena, this means that subpersonal explanations can be constitutive explanations. So, the plan is to argue first for the constitutive relevance of the subpersonal level via a limited application of the isomorphism principle, then for the availability of in some cases of non-isomorphic explanations – giving reasons to think that at least some perceptual experiences fall into the category. These, *perspectival* experiences are those that situate their subjects in relation to objects in the environment that are causally responsible for perceptions and causally responsive to agency.

In section (§2.1), I outline and motivate the claim that psychological phenomena constitutively depend on conditions under which subjects are interpretable as manifesting them; I argue that

[13] The continuity of this with later 'mature' versions of the subpersonal/personal distinction in Dennett's work is itself matter of interpretive debate (see, e.g., McDowell 1994; Hornsby 2000; Drayson 2012, 2014; Colombo 2013; Rupert 2018). My reading of Dennett diverges from the strong claim, widely assumed – and defended by Hornsby (2000) – that Dennett's early view is radically different from the later picture and that the former is more hospitable to 'realism' about personal level phenomena. I think that Dennett's interpretationism stakes out a position that is difficult to fit within the realism/anti-realism dialectic, and that this aspect (inherited in a modified form from the philosophers of the personal level) is a persistent in his developing picture of the personal/subpersonal distinction. For more discussion, see (§2.2).

Dennett introduces the distinction between personal and subpersonal levels to show that vertical explanation in terms of neurocomputational can supply criteria of interpretability. In (§2.2), I spell out the notion of horizontal explanation in some more detail, distinguishing between causal and non-causal models. Here, I focus on the non-causal models proposed by Ryle and Anscombe, whom Dennett identifies (along with the later Wittgenstein) as the paradigmatic philosophers of the personal level. In (§2.3), I introduce vertical interpretation via a reconstruction of Dennett's original distinction between explanatory levels. This supports a model of subpersonal explanation as content-ascriptive mechanistic explanation of personal-level phenomena. In (§2.4), I make this account more precise, summarising a widely-accepted account of how subpersonal explanations connect psychological phenomena with neurocomputational mechanisms and introducing the explanatory isomorphism constraint. Finally, in (§2.5), I contrast subpersonal explanations of migraine aura and ordinary visual object perception. I note that while explanatory isomorphisms do hold in the aura case, they do not for object perception. Reflecting on differences between these cases supports non-isomorphic models of subpersonal explanation; and *these* suggest that episodes of perspectival experience are vertically constituted by neurocomputational mechanism's representational states and horizontally constituted by persons' extra-mechanistic relations to their environmental contexts.

1) Psychological Interpretation and the Personal/Subpersonal Distinction

a) Subpersonal mechanisms as 'enabling conditions'

The distinction, first explicitly introduced under this label by Dennett (1969), between "the explanatory level of people and their sensations and activities" and "the sub-personal level of brains and events in the nervous system" (p. 93) is, as we've already seen above, sometimes appealed to by defenders of relational or naïve realist models of perceptual experience. For example, McDowell has suggested that we should understand the distinction as "a special case of a more general distinction between content-attribution at the level of the animal and content-attribution at the level of its internal machinery" (1994, p. 201). His point here is that the philosophical questions that relationalists are trying to answer target the content and phenomenal character of states of the whole animal – the subject of experience – and *not* its internal machinery. Content attribution at the personal level is supposed to be a fundamentally different kind of project to content attribution at the level of mechanisms, one which considers the subject holistically against the backdrop of her rational capacities and psychological history, tracking the real content of states and attitudes that connect her to her actual environment. By contrast, content attribution at 'lower levels' is an explanatory heuristic, derived from and

parasitic on this ‘genuine’ personal-level content attribution. Since the mechanisms don’t meet criteria by which psychological content is attributed, their informational states should not be understood as having genuinely psychological content:

> At the level of internal machinery it is useful to talk of sensory systems as information-processing devices but for the animal its sensory systems are modes of openness to features of its environment. Information-processing characterizations of the internal machinery figure in explanations of how it can be that animals are in touch with their environments. The 'as if’ content that is usefully deployed at the lower level helps make intelligible the genuine content that appears at the higher level by way of ‘enabling explanations’ not as somehow constituting that content. (pp. 201-202)

I take it that McDowell’s view is something like this: While content attributions at the level of mechanisms ‘make intelligible’ how there could be genuine content at the personal level, what this amounts to is just a way of giving convenient labels to mechanistic activities and processes. While these ‘information-processing characterisations’ might be useful to researchers trying to individuate mechanisms or predict behaviour, they are not available to the subject for conscious reflection. Subpersonal contents do not enter into the ‘space of reasons’ (Sellars 1956; McDowell 1996) occupied by the contents of her perceptually grounded beliefs and intentions, which draw from the genuine content of personal-level experiential states. Neurocomputational mechanisms and their representational content underlie and support, but do not constitute, a subject’s capacity for experience.

It is important here to distinguish claims about identity from claims about constitution. The negative claim that the relationship between psychological and neurocomputational phenomena is *something other than identity*, is widely endorsed by philosophers of mind and cognitive science – often for reasons that have nothing to do with the motivations behind relational views like naïve realism. There are consequently many ways that we can make sense of neurocomputational explanations of perceptual phenomena without identifying personal-level phenomena and the neurocomputational mechanisms evoked to explain them. French and Philips (2023, p. 365) might seem to be making this weaker non-identity claim, when they argue that

> the naïve realist can endorse approaches to perception on which the visual system is seen as computing distal causes from proximal inputs, with the sole proviso that

> the products of the relevant inferences are not construed as representational *states identifiable with perceptual experience.*

But since French and Philip's proviso is met by Beck's (2019) neurocomputational naïve realism, we should by now be wary of taking it at face value. And as they flesh out the point, it becomes clear that a more demanding condition than non-identity must be met. As in McDowell's (1994) presentation, the promise of the subpersonal/personal distinction is supposed to be that it allows us to distinguish sharply between "the question of what perceptual consciousness consists in", on one hand, and "questions about the neurological and psychological causes of such experience, and the processes underlying it" (French & Philips 2023, p. 365), on the other. Since ascribing a relation of constitution is a weaker metaphysical claim than ascribing identity, ruling it out excludes more than ruling out identity.

To give a standard example, Michelangelo's *David* is made up of – *materially* constituted by – a particular physical mass of marble. But many philosophers have resisted the claim that the *David* and this marble are identical (see e.g., Baker 1997).[14] If identity were all we needed to rule out, we might think that we could understand the constitutive role of perceptual mechanism in a similar way – we could appeal to something like the notion of material constitution to make sense of the relationship between personal and subpersonal levels (cf. Bechtel 2008; Yilkoski 2013). But this isn't what French and Philips have in mind. Rather a sharp distinction between enabling causation and constitution is key (for naive realists) because it:

> opens up the possibility that the images or representations posited by perception scientists are part of the story concerning the processing which causes and underlies perceptual experience, but *no part of the answer to the question of what experience consists in.* (ibid., emphasis added)

On this vision of the personal/subpersonal distinction, the condition on naïve realists endorsing neurocomputational explanations in cognitive science is that *only* the second – causal – kind of question pertains to neurocomputational processes and mechanisms (and vice versa). If naïve realists want to capitalise on the possibility opened up by distinguishing between the constitutive questions asked by philosophers and the causal questions of perception scientists,

[14] One reason for this is that the marble has different modal properties to the statue. It might survive events that everyone agrees would destroy the statue. It (or much of it) could have formed a different statute – as, in fact, it nearly did on two occasions before Michelangelo received his commission. Many have also claimed that the *David* has properties – most saliently, various kinds of aesthetic and social value – that don't attach to the marble it is made of. Both lines of thinking are closely related to what I'm calling horizontal constitution (see §2.2a).

they should assign neurocomputational mechanisms no part in their *constitutive* personal-level accounts. This, however, is just what Beck's proposal does. While neurocomputational appearances are not identical with the phenomenal character of perceptual states, they are supposed to be an element of what makes them up, since they figure as relata in the perceptual acquaintance relation.

This reading of the personal/subpersonal distinction[15] gives us a neat diagnosis for the problems highlighted by the previous chapter's discussion of Beck's neurocomputational naïve realism: freely crossing distinct metaphysical levels to construct a model appropriate to phenomena at only one level, it adopts the critic of naïve realism's assumptions about the proper constituents of experience and so deprives its main claim of the motivation and support it could take from the distinction (cf. Epstein 2022, p. 4 ft. 16). But getting *this* kind of motivation and support for naïve realism requires us to adopt a reading that its critics can simply reject. French and Philips might be right in characterising non-identity between personal-level and subpersonal representations as a relatively un-burdensome commitment for naïve realists to take on (I think they certainly are right to present it as a necessary one). But claiming that representational mechanisms play only a causal or enabling role, or that contents ascribed to subpersonal mechanisms have a derived, heuristic status, means staking out strong, contentious positions in the philosophy of cognitive science. To adopt them, they incur costs that opponents of naïve realism are unwilling to pay, shifting the ground of discussion from specific considerations about the nature of perception and its relation to underlying mechanisms onto more general considerations about naturalism, ontology and explanatory reduction.

Ironically, debates of this kind were what Dennett's original proposal of the personal/subpersonal distinction was meant to avoid (Dennett 1969, pp. 1-6). Counting on a connection between the explanatory and constitutive relevance of psychological interpretation that was already in currency at the time of its publication, the model of subpersonal explanation proposed in *Content and Consciousness* (see also Dennett 1968, 1970) aimed to legitimise mechanistic explanation in the then-emerging field of cognitive neuroscience as genuinely psychological, and *therefore* potentially constitutive, explanations (cf. Davies 2000; Drayson 2012). What's more, the notion of non-isomorphic subpersonal explanation set out in his (1978, 1989, 1990) and Hurley's (1998) later developments of the idea, offers a promising basis for

[15] French and Philips prefer a distinction between 'ordinary' and 'psychological' explanations (p. 372, ft. 14; cf. Philips 2019) that won't track the personal/subpersonal distinction on all construals – they remain 'neutral' about this (ibid.) But since *all* subpersonal explanations are 'psychological' explanations in their sense, we can treat them as equivalent for present purposes.

accommodating the central naïve realist claim – that objects of experience partly constitute experiences of them – without denying the constitutive relevance of neurocomputational mechanisms. The idea, in brief, is that for psychological phenomena 'genuine' constitutive explanations are enabling explanations – that being able to say how a capacity or state is realised gives you all you need to be able to say what it is. Given these potential advantages, this model deserves more sympathetic attention from naïve realists than it has typically received.[16]

b) Subpersonal explanation and vertical explanation

Naïve realists are not alone in projecting their prior metaphysical commitments onto the personal/subpersonal distinction. The labels 'personal' and 'subpersonal' have been attached to a wide range of entities. Reviewing the 'uses and abuses' to which the distinction has been put, Zoe Drayson (2012, p. 8) lists talk of personal- or subpersonal-level "events, processes, states, contents, mechanisms, and facts." And realism about at least some of these putative entities is embedded in what Rob Rupert (2018) has described as the 'received view' of the distinction. According to Rupert, the received view involves three distinctive commitments: a) that psychological phenomena (beliefs, perceptions, intentions, etc.) correspond to entities at a distinct and unified 'level of reality'[17]; b) that the personal level is a domain of knowledge that is prior to and independent of scientific investigation; c) that cognitive scientific explanations target psychological phenomena and specify their constitutive mechanisms. Let's call (a) the *ontological claim*, (b) the *autonomy claim* and (c) the *methodological claim*.

Rupert thinks that we should reject all three claims and consequently that we should abandon the notion of explanatory levels introduced by the personal/subpersonal distinction. But while I think it's fair to say that all of the claims that make up the received view are widely held, it's worth questioning whether they are (or should be) held *together*. On the face of it, each can be evaluated independently of the others. We've already considered views on which they seem to come apart: French and Philips endorse an *ontological* distinction between levels, and McDowell seems to be arguing towards this from some version of the *autonomy* claim, but neither endorse the *methodological* claim – rather, it looks like this is part of what they're

[16] A recent *favourable* discussion of Dennett's early work within the naïve realist framework comes, tellingly, in Ward's (2023) proposal of a hybrid enactivist-naïve realist approach.

[17] I take it that there being levels in this sense entails, minimally, that entities at different levels cannot be identical.

arguing *against*. On the other hand, we have little reason to think that Dennett takes on the ontological commitments of the received view.[18]

As Drayson (2012, 2014) convincingly argues, the core distinction for Dennett (1969, 1978) is between personal and subpersonal *explanations* rather than personal and subpersonal *entities*. We can better understand the role it plays in this work and the influence it is had on the subsequent development of philosophy and cognitive science when we understand that the notion of subpersonal explanation is supposed to introduce a *vertical* dimension to genuinely psychological explanation (Bermudez 2000; Drayson 2012). Subpersonal explanations are, for Dennett, psychological (they appeal to psychological explanatory norms) and at least sometimes vertical (they cross the explanatory levels of persons and mechanisms).

But what would it mean for psychological explanation to have a vertical dimension? Contrasts between vertical and horizontal explanations are a common feature of many debates in metaphysics and the philosophy of science (see, e.g., Kim 1993; Bechtel & Abrahamsen 2005). Drawing on these, it might be helpful to introduce the two directions of explanatory dependence in general terms at first:

> *Horizontal explanation:* An explanandum, A, is horizontally explained by an explanans, B, iff A and B are temporally ordered (standardly, A follows B) and A depends on B according to some relevant principle of explanation.

Sets of horizontal explanations determine *levels* of explanation. That is, for an explanatory system consisting of multiple levels, any one level is determined by all the elements of that system that can be understood as interacting in a way tracked by a horizontal explanation. On this way of understanding the notion of an explanatory level, levels of explanations come after explanations themselves. Explanations are not thought to hook onto preexisting 'levels of reality' but rather the notion of an explanatory level is an idealisation that emerges from the ways that successful explanations fit together. Thus, weather fronts, economic depressions and the beliefs of individual persons exist on different levels of explanation not because there exist discrete levels of reality (or scale, or complexity) that correspond to weather, economics and individual psychology. Of course, as these examples suggest, there are differences in degrees of complexity or scale that are associated with explanatory levels – and these differences are part of what accounts for the explanations offered by the discourses that apply to them patterning

[18] It's unclear in Rupert's discussion whether Dennett is supposed to be an advocate of the received view. Though Rupert criticises aspects of the distinction as presented in Dennett's work, he does not include Dennett on a list of "influential philosophers of mind" who hold it (Rupert 2018, p. 7).

as they do. But we needn't assume a systematic connection between something as abstract as scale or complexity and the levels that we in fact use. Sometimes complex things are affected by simple things, sometimes small things interact with large things. Where there are reliable explanatory principles that ground a scientific or otherwise dependable explanatory discourse that account for these interactions and effects, these principles are part of what determines the explanatory level on which these are accounted for. Whatever they are, where A and B are linked by an explanation of the kind just outlined, we can say they are on the same explanatory level.

Not all targets of explanation are related in this way. Material constituents of things, for example, do not horizontally explain the properties of the things they constitute. When we appeal to something's material constitution to explain its properties, we rely on a different kind of relation. These do not hold across time but between levels of scale or complexity *at* a time (or over an interval).[19] These vertically relate two (horizontal) levels:

> *Vertical explanation:* A is vertically explained by B iff B is at a different level of explanation to A (standardly, B is at a 'lower' level than A) and A depends on B according to some relevant principle of explanation.

The horizontal/vertical distinction is often associated with the distinction between causal and constitutive explanations. This association holds when, for example, material properties of an object are used to explain the changes it undergoes as a result of some causal interaction. A ball thrown at a window *horizontally* explains its having broken (the event of the ball's being thrown precedes and causes the breaking of the window) – and we provide an additional *vertical* explanation of the latter event when we appeal to the window's material constitution (and so, its fragility) as part of an explanation of why *this* cause had *that* effect. But, as we'll see below (§2.2), not all constitutive explanations are vertical. Nor is it obvious that all vertical explanations are constitutive or that all horizontal explanations are causal.

Paradigm examples of psychological explanation are horizontal.[20] For now, we can characterise psychological explanations as targeting persons, their activities, states and behaviour. An explanation is psychological in the relevant sense if a psychological phenomenon figures either

[19] A common view about *constitutive* vertical explanation is that it relies on synchronic interlevel relations that hold in the same way at every instant that the explanatory relation holds (see, e.g., Yilkoski 2013). This view is closely related to the explanatory isomorphism principle that is denied application to perspectival experience below (§2.5). Though the justification for this will have to wait until then, it should therefore not be assumed as part of a neutral characterisation of vertical explanation.

[20] Note that this can be true whether or not ordinary psychological explanation is best understood as causal explanation.

as explanans or explananda and the explanatory principles involved are distinctively psychological. Paradigmatic psychological explanans/explananda are characterised in intentional terms (as having intensionally opaque content), and their explanatory relations reflect psychological norms that specify how the different elements of a subject's psychological life and behaviour should relate to each other – in accordance with laws, explanatory generalisations, or idealised patterns manifested in the thought and behaviour of psychological subjects. 'Rationalising' explanations (Davidson 1980) citing an agent's reason to act in a certain way fit both of these criteria – and consequently so do many that cite a perceiver's experiencing an object or event relevant to her reasons for acting. Important psychological explanatory principles include norms of holism (subjects tend to take all of the reasons available into account) and consistency (where two activities, ϕ-ing and ψ-ing, are inconsistent, reasons to ϕ tend to explain a person's ϕ-ing and not her ψ-ing), and norms of environmental appropriateness (the environment in which S is actually located sets standards of success for her activity).[21]

What about neurocomputational explanations? These also attribute intentional content to mechanisms and explain their activity by appeal to norms of consistency and appropriateness. Interpreting a perceptual mechanism as drawing inferences about an animal's environment or representing some feature or property in the environment grants explanatory purchase on the mechanism via distinctively psychological or ideally rational norms – this is what McDowell has in mind when he describes neurocomputational explanation as 'heuristic' or the attribution of 'as if' content. Such heuristics explain *on the model of* personal-level explanation. Cognitive scientists treat mechanisms as being 'at the same level' when they offer explanations of this kind, for example explaining stages of perceptual processing as computational steps implemented by causally interacting mechanisms: such explanations are horizontal at 'one level down' from paradigm personal-level explanations.[22] But not all neurocomputational

[21] For more precise characterisation of these norms and discussion of their relations to each other, see (§§2.1c, 2.2b)

[22] Some philosophers (e.g. Rupert 2018) argue that this lower level is in fact all that psychology needs. Supporting this conclusion, however, requires us to extrapolate from the success of horizontal neurocomputational explanation on comparatively short timescales and in controlled condition to assign it a comparable level of prospective success to that enjoyed by ordinary interpersonal interpretation and explanation over much longer timeframes and in a much wider range of circumstances. That is, it requires us to go well beyond the explanatory practices we have, based on the metaphysical argument that since the mechanisms realise the capacities targeted by ordinary explanation, this kind of explanation will ultimately reduce to explanation cast in terms of mechanism. Not only does this make questionable assumptions about what realizes the capacity (the mechanisms, rather than the mechanisms in their worldly context) but it also mixes up the metaphysical notion of a 'level of reality' with the methodological one of an explanatory level. It should be noted that there can be levels of

explanations are horizontal. Implementation of neurocomputational properties is explained *vertically:* neurophysiological causal relations are at a lower level of explanation to the computations they implement. But representations, too, can be understood as underlying patterns of horizontal explanation at the 'higher' psychological, levels. Interlevel explanatory connections are involved when researchers assign psychological roles to structures and activity in the brain, via neurocomputational contents that realise a possible implementation of that psychological phenomenon.[23]

This picture, which shaped the mainstream of research after the 'cognitive revolution' in the mind sciences (Cummins 2000; Miller 2003), is subject to philosophical objections. And it was primarily these objections that Dennett (1969) had in view when he introduced the subpersonal/personal distinction – the book was intended as a manifesto for cognitivism in philosophy and psychology (Dennett 2008; Ross 2015). Importantly, these objections are not grounded in the ontological considerations tabled by McDowell, French and Philips. But the waters have been muddied by two factors: a) Dennett's wider project understands the constitutive conditions of psychological (and neurocomputational) states in terms of the same set of explanatory norms that many other philosophers have used to distinguish *between* personal and subpersonal levels; and, b) this approach, which I'll call 'interpretationism' (Schwitzgebel 2006; Pautz 2021b), does not commit us to a ruling one way or another about explanatory or metaphysical priorities between levels. As I understand it, Dennett's overall picture of the relation between personal and subpersonal levels rejects the *ontological* but endorses the *methodological* claim, while introducing a modified version of *autonomy*, which I'll call the *explanatory interdependence claim*. This denies independence of the personal level for content attribution while accepting its priority. While the personal level is where interpretation starts, personal-level content attributions can be revised, and explanatory norms contravened, for reasons revealed by subpersonal explanations.

c) Subpersonal explanation and the interpretationist framework

Interpretationist approaches in the philosophy of mind give a central role to practices of psychological explanation that people use to interpret and predict each other (and, sometimes,

---

explanation, in the latter sense, even if there are no levels of reality. So, even granted the success of visionary arguments due to Rupert and others against the metaphysical distinction between a person and her psychological capacities on one hand and the action and operations of mechanisms that realise those capacities on the other, it is unclear that an argument based on levels of explanation will be affected (especially if, as I suggest, identifying a *neurocomputational* mechanism level requires that we retain its distinctness from the personal level).

[23] This rough picture will be sharpened in (§2.3)

themselves). The role given to these practices, however, differs widely across interpretationist accounts. Two opposite ends of the interpretationist spectrum are occupied by Lewis (1973) and Anscombe (1957). For Lewis, everyday psychological interpretation anchors an ideal, formal theory by which content ascriptions could be derived from non-psychologically charaterised sets of basic fact. On this view 'interpretation' "is not any real-life task of finding out about [a subject's] beliefs, desires, and meanings. [It does not ask] how we could determine these facts. Rather: how do the [non-psychological] facts determine these facts?" (p. 333; cf. Williams 2020). Anscombe, on the other hand, resists the reductionist aims embodied in this picture; what one learns about the structure of agency on her account, is learnt as a participant, from within the patterns and practices of everyday psychological interpretation. Nonetheless, in setting out the principles guiding this practice we do not "describe actual mental processes" but rather "an order which is there wherever actions are done with intentions" (Anscombe 1957, §42). On this spectrum, I think that both Dennett and Hurley both fit closer to the middle than to either extreme pole: interpretability encompasses the physical facts, but these are accessed through cognitive scientific explanation that begins with the ordinary interpretive task of discerning a psychological order that is already there in them.

A motivation uniting these proposals is that reflecting on practices of psychological explanation can generate solutions to philosophical problems that look intractable when detached from these practices. For example, consider a radically anti-interpretationist *Cartesian view* of psychological phenomena.[24] The *Cartesian view* identifies these with a person's internal states and events and holds that their natures are determined by their intrinsic properties – a person's psychological interpretability is a contingent outcome of the way her inner states direct her movement and talk. There are many philosophical worries about this picture. As we saw in the previous chapter (§1.1a), one group of these centre on several kinds of alienation (e.g. epistemic, semantic, agential) that the subject so-conceived stands in with respect to the world. The acquaintance-theoretic response to these is to introduce a distinctive kind of psychological relation, so that the nature of (at least some) psychological phenomena is not determined by the intrinsic properties of states or events internal to the subject. The interpretationist response is to more generally reorient our understanding of psychology from inner entities that are constitutively independent of each other and the world to those patterns of interaction (with each other and, through perception and agency, with the environment) that the *Cartesian view* presents as their contingent inputs and outputs. As with the acquaintance model, this shift

[24] In providing this caricature position, I don't mean to attribute the claims that make up the *Cartesian view* to Descartes or the early modern Cartesian tradition.

emphasises relations in which subjects stand rather than intrinsic properties of their inner organisation. But these relations are neither basic nor (necessarily) introspectively transparent. The patterns of relation between subject and world that interpretationists focus on are individuated by the psychological interpretations suitable to them, and these place in them a network of other psychological phenomena. It is essential to any token psychological phenomenon that it is attributed to an interpretable subject, but being an interpretable subject requires that one's psychological states and capacities form an interconnected system.[25]

Interpretationists think that a minimal conditions on the claim that, for example, a subject, S, believes that *p* is that she be interpretable as believing that *p*. Here, 'interpretability' corresponds to a set of constraints on the relations between psychological phenomena and behaviour that allow interpreters to determine what S believes. A similar condition on perception ties someone's being in a perceptual state with her acting and forming beliefs in a way that is responsive to the environment as she perceives it – and so, for example, *reporting* things to be as her perception presents them as being. In as far as these downstream events and activities bear on interpretation, they help her meet the standards of interpretability that are at the centre of the account. I'll call these I-criteria.

I-criteria come in both necessary and sufficient versions:

> *I-criterion (necessary)* : If a subject, S, is in a psychological state Ψ, then S will be in a condition C^{Ψ} such that knowing that S is C^{Ψ} justifies the interpretation <S is Ψ>
>
> *I-criterion (sufficient)* : If S is in C^{Ψ} then S is Ψ

While interpretationism has sometimes been understood as entailing both of these claims, I take it that only the necessity claim is essential to the view (cf. Davidson 1984; Child 1994 pp. 23-40). Tying conditions of interpretability to the nature of psychological states is consistent with interpretation being fallible and with the possibility of interpreters being justified but wrong. What the connection rules out is that someone could be in a psychological state while being systematically uninterpretable as being in that state. In Davidson's formulation, interpretationism entails that what "there is to learn" about a person's intentions, beliefs and desires coincides with "what a fully informed interpreter could learn" (1984, p. 315). Strong

[25] A more radical interpretationist proposal links constitutive interpretability with the actual presence of an interpreter. On this view, it's a condition for having a mind not just that one be interpretable, but that one is at least sometimes actually interpreted. In this case, the mental phenomena of other subjects (i.e., interpreters) are partly constitutive of psychology of the target of interpretation. I will not defend or presuppose this radical view (for discussion, see Williams 2020, p. xix)

versions of interpretationism appeal to I-criteria that set relatively strict limits on what is accessible to fully informed interpreters, while weaker versions allow more information to the fully informed (and increasingly ideal) interpreter.[26] These will correspondingly count a wider or narrower range of things as determining I-criteria.

Much of the philosophical significance of interpretationism therefore depends on how conditions of interpretability are understood – questions about what kind of thing, state of affairs or relation gets to be a value for C^{Ψ}. Before this is decided on, many possible consequences of interpretationism for debates about reduction and realism remain open. So, as I'll understand it, 'interpretationism' refers to a way of answering philosophical questions about the mind based in reflecting on the nature of psychological explanation, rather than a metaphysical theory of psychological phenomena that might be opposed to other standard views.[27] Claims about realism, naturalism and interpretation can vary independently. So, for example, one might characterise identity conditions of 'robustly real' psychological phenomena in terms of their contribution to interpretability; or one might claim there is 'nothing to' these phenomena beyond the interest- or observer-relative fact that it's useful to interpret persons in terms of psychological phenomena. Both views are interpretationist, but only one is realist.[28]

However interpretability is understood, finding out whether someone or something meets I-criteria additionally involves assessing that person or system in light of the psychologically and rationally normative standards mentioned above. For example, attributing the belief that *p* to S requires this content to be checked for consistency against the content of other attributed beliefs. While some inconsistency is tolerable, it comes at an interpretive cost. The more inconsistency, the less explanatory purchase the attribution of any one content will give an interpreter on S's behaviour. Since, in general, attributing the belief that *p* to S is can be explanatory only insofar as it rules out the belief that ~*p* (and any behaviour consequent on it),

[26] The limit case, an omniscient interpreter, will set no constraints on interpretability conditions (for her, since she knows everything, the psychological states of any subject will be self-disclosing). In a Cartesian world (in which the nature of psychological phenomena is fixed by intrinsic properties that only contingently supply public interpretability conditions), the omniscient interpreter's assignment of psychological conditions would be extensionally equivalent to those predicted by the *Cartesian view.*

[27] Child (1994) defends a similar conception of the interpretationist project. There is plenty of evidence that Dennett thinks along similar lines (in this respect), for example where he sets out the goal of X as Y. Williams (2020) demurs, casting interpretationism as just such as substantial metaphysical programme – albeit, one in need of supplementation from realist causal theories. (see also Lewis 1973).

[28] As more than one commentator has noted (see Drayson 2014; Shea 2018, p. 13-14) it can be hard to pin down Dennett's stance on classic metaphysical questions about the mind. But on one view often attributed to him, he endorses the second approach. This view, call it 'instrumentalism', combines an interpretationist approach to identifying psychological states with further anti-realist claims in metaphysics. For more discussion, see Child (1994, ch. 1), Williams (2020, p. xix)

inconsistency undermines explanation by weakening the constraints that belief attribution puts on possible interpretations. If these constraints are too weak, S will not be interpretable in terms of belief and so fails to meet I-criteria for having beliefs.

Anything that meets I-criteria can be understood and predicted in a way that requires very little (explicit) understanding of its inner causal dynamics. Consequently, many interpretationists think we can appeal to psychological explanatory norms in constitutive explanations – explanations of the basic psychological attributes of persons – that do not involve a level of explanation lower that the interrelated states themselves. Some have concluded from this that interpretationism is incompatible with the received view's *methodological claim* because it suggests that psychological contents will (at least) depend on the condition of the interpreted system as a whole – including the relations in which it stands to its environment. The content of an intentional state is essential to its identity. So, if assigning content depends on the condition of the whole cognitive system, it looks like the intentional states can't be reduced to the activity of anything less than the whole system. Attempts to do so, on one influential way of developing this thought, commit the 'mereological fallacy' – they explain a property or capacity realised by the system as a whole as though it only depended on one of its parts (Hacker & Bennett 2003)

This argument has roots in the philosophers of the personal level. Wittgenstein's discussion of psychology can be seen as treating some types of behaviour as I-criteria, for example in the claim that "only for a being that behaves like a human we can say that it has pain" (1953, § 283). Ryle's (1949) notion of a category mistake involves the misattribution of systemic properties to parts of systems – and is also introduced in the context of a discussion of the metaphysical commitments of talk about the mind. The "special sense of the question 'why?'" at the centre of Anscombe's (1957) treatment of intentional action applies to the behaviour of agents – understood as integrated loci of practical rationality fitted into a context only specifiable in intentional terms. In as far as these arguments play a negative role, undermining traditional "mechanistic and paramechanistic hypotheses" (Ryle 1949) about mental entities, they are also crucial to Dennett's (1969) formulation of the personal/subpersonal distinction. But to see why, for Dennett, these negative insights are compatible with retaining the *methodological claim* and replacing *autonomy* with *interdependence* we should first get clearer on their associated positive programme.

2) Interpretation and the Philosophers of the Personal Level

a) Interpretability and horizontal constitution

Since the negative case made by the philosophers of the personal level rules out traditional approaches to identifying psychological phenomena, they need some other way to do this. I want to suggest that a distinctive kind of *horizontal* constitution plays this role. First, here's a general characterisation:

> *Horizontal constitution:* A condition, C, horizontally constitutes a psychological phenomenon, Ψ, iff Ψ depends for its identity on actual or possible relations between Ψ (or its derivatives, ψ_1, …, ψ_n) and C (or its components, c_1, …, c_n)[29]

Not everything has horizontal constitution conditions. Horizontal constitution accounts for the identity of phenomena that depend for this identity on interactions or relations with their contexts. Paradigmatically, these include 'social kinds' (Ásta 2017) such as money or citizenship (and consequently the property or state of *being* money or *being* a citizen). They likely include more than social kinds, and they might include most kinds (for discussion, see Hacking 1991; Dupré 2017; Ásta 2017). Introducing the category through social kinds, however, usefully brings into focus an important parallel between these and psychological states. Both are at the centre of practices of explanation of the behaviour of psychological agents that can be highly articulated and predictively powerful without (necessarily) quantifying over entities at any lower level of description.

Successfully engaging in these practices needn't – and often doesn't – involve finding referents for their terms in a more fundamental description of the objects and processes involved. In the context of the debate about naïve realism, Campbell (2020) illustrates a similar point with an analogy between what he takes to be the explanatory autonomy of 'high-level' perceptible properties like objective colour and the autonomy of the high-level explanatory notion of a 'credit crunch' from explanations couched in terms of fundamental physics:

> So far as the high-level economic properties go, there are quantitative models for what is happening: the failure of confidence, the liquidity crises for businesses, and so on. […] But that's not to say that we should expect there to be quantum-mechanical structures, with independently identifiable quantum-mechanical

[29]The parenthetical clauses about derivatives and components are meant to cover cases where 'relations' are interpreted as actual causation (cf. Davidson 1980). In this case a psychological state might be thought to cause or be caused by a change in (a component of) a wider condition via a derivative mental event (e.g. an 'onset' of the psychological state).

> significance, that correspond to the structure of a credit crunch. (Campbell 2020, p. 414)

A possible further analogy between economic and psychological phenomena, not discussed by Campbell, is that the relevant practices of explanation are closely tied to the constitutive conditions themselves. If people could not describe and explain interactions involving, for example, the exchange of money in terms of the economic values involved, then it is far from clear that these interactions would constitute this kind of exchange. Quantitative economic models, like those mentioned by Campbell, are explanatory in as far as they abstractly and generally describe what is done agents that regulate their concrete behaviour partly in terms of similar abstractions (contrast: evolutionary models in biology).[30] The philosophers of the personal level, as well as later interpretationists like Dennett and Davidson, exploit both parallels, finding important connections between the constitutive conditions of psychological states and the practice of ordinary psychological explanation. Being an appropriate target for horizontal explanations of a certain kind is part of what it is for things to be with you, psychologically, as they are.

We might think that psychological phenomena are horizontally constituted by their explanatorily relevant contexts because ordinary psychological explanations are horizontal: they don't account for what they're trying to explain in terms of underlying entities but rather in terms of relations – between psychological phenomena themselves and between these and the environment. In this case, horizontal constitution conditions are given by characterising the appropriate kinds of explanation and the relations they involve. The next step is to say how explanations of this kind work.

b) Non-causal explanation at the personal level

A widespread understanding of personal-level explanation sees it as causal: psychological phenomena are particulars and their interface with the world is mediated by causal interaction between psychologically and non-psychologically characterised events. One version of this position, call it the event-causal model of explanation, is defended by Davidson (1980) – and this is at the root of what has come to be known as the 'standard story' of action explanation (Aguilar & Buckareff 2010). The standard story and its relation to the subpersonal explanation will be the focus of much of the next chapter's discussion (§3.2), so I'll only give a brief outline here:

[30] Note: the claim that economic agents (as modelled) act for economic reasons is distinct from the claim that economic models derive from *psychological* agents' reason-guided behaviour. In the text, I intend to make only the former, less controversial, claim (for discussion, see Ross et al. 2012).

Explanations of why something happened can be psychological, according to the standard story, if psychological states and events figure – in the right way – in the causal chain that led up to that thing happening. What counts as the 'right way' is determined by psychological norms, to which the causal chain has to conform. On this view, an explanation's psychological character attaches to the properties of the causally related phenomena involved and the way these are reflected in patterns of event-causation.

While Davidson's (1980) and later interpretationists' accounts (see, e.g., Lewis 1973; Williams, 2020), build around this picture, it is rejected by the philosophers of the personal level.[31] For example, consider this (in)famous passage from Ryle (1949, p. 87), in which he contrasts two possible construals of the statement 'he boasted from vanity'. Respectively, these suggest radically different ways of understanding how the psychological condition of the person mentioned (his vanity) explains a corresponding piece of behaviour (his act of boasting):

> The statement […] ought, on one view, to be construed as saying that "he boasted and the cause of his boasting was the occurrence in him of a particular feeling or impulse of vanity"

While, on the other view,

> it is to be construed as saying that "he boasted […] and his doing so satisfied the law-like proposition that whenever he finds a chance of securing the admiration and envy of others, he does whatever he thinks will produce this admiration and envy"

The explanandum is the same on both construals (*that* he boasted when he did – a fact about an agent grounded in a particular event), but the two proposed explanations differ in their structure and, correspondingly, in their ontological commitments. The first explains one singular, dated event in the shared social environment as the effect of an antecedent (singular, dated) psychological event that is in itself private to its subject. The causal antecedent is identified independently[32] as a feeling or impulse of a certain kind, possibly characterised by its subjectively felt character ('a particular feeling'). This neatly fits the causal model of horizontal

[31] As Child (1994) notes, there are reasons to think that the causalist picture and interpretationism are in tension. He offers one way to resolve this tension – and, as will become clear in the following chapter, the two-level interdependence view offers another. Having got more of the alternative view on the table, this chapter (§§3.3-3.4) will return to Davidson's causal model in more detail.

[32] This allows it to fulfil the Humean condition on informative causal explanation that causes and effects be independently identifiable (Lewis 1974). If the second explanation is interpreted as an event-causal account (e.g., as the claim that boasting is caused by boastfulness'), it will not meet this condition.

explanation; it introduces events and relates them causally over time in a way that is at least roughly analogous to ordinary event-causal explanations involving no psychological entities. But this comes at the cost of attributing intrinsic properties to the inner psychological causes of behaviour that give rise to the philosophical worries canvased above. According to a popular view of the contingency of causal relations, often associated with Hume, it rules out horizontal constitution: if cause and effect must be independently identified, then a psychological phenomenon must be constitutively independent from associated behaviour. Identifying features must be sought elsewhere – with the introspectable intrinsic properties of the traditional picture being a natural place to look.[33]

The second construal avoids explicit reference to mental events independently characterised in terms of their intrinsic properties. Of course, Ryle's paraphrase talks about the boaster's 'finding' opportunities to secure admiration and envy, 'thinking' and 'acting' in corresponding ways, and even "indulging in roseate daydreams about his own successes" (p. 86), but these are used to characterise a psychological disposition; and it's this disposition that carries the burden of explanation. The fact of the boaster's boasting is not explained by reference to an *act* of boasting and a particular psychological cause. Rather, the statement 'he boasted from vanity' subsumes the whole episode under an explanatory generalisation (a 'law-like proposition')[34] characterised in terms of the psychological disposition manifested in the behaviour.

Ryle argues that we should put dispositions, rather than events, at the centre of our understanding of psychological explanation. The payoff of this is supposed to be that interpretation and explanation of other psychological agents need not involve mechanistic or paramechanistic hypotheses about their inner workings (and that self-understanding need not involve introspective access to this machinery). This suggestion underlies the negative claim that psychological terms don't refer to special 'psychological' causes of behaviour and, collaterally, that psychological explanations don't reduce to event-causal explanations at a lower level of description. Respectively, these denials rule out the 'paramechanistic' and 'mechanistic' options, as Ryle sees them.

The lesson can be generalised from the use Ryle makes of this approach in his disposition-based proposal. For example, Anscombe's (1957) treatment of intentional action explicitly rejects a

[33] They aren't, of course, the only place to look. Isomorphic *subpersonal* explanations introduce other, naturalistically motivated candidates. But it's a notable feature of such accounts that they often *start* from considerations grounded in introspection (see §3.1a below)

[34] Ryle's invocation of explanatory 'laws' here accounts for much of this passage's infamy. But there are good reasons to read Ryle's usage as very loose (cf. Dennett 2000). Correspondingly, I intend the notion of an explanatory generalisation to be read broadly.

dispositional account of action or intention (see Anscombe 1957, § 13), but makes three interdependent claims that together distinguish it from the event-causal model:

a) Movements of an agent's body can be intentional only under a description; a movement that is intentional under one description can be unintentional under another (§ 6)
b) Intentional actions cannot be analysed as a bodily movement and a psychological cause of that movement (§§ 10-16)
c) Intentional actions are intelligible in light of reasons. Reasons for actions are given in answers to 'a certain sense' of the question 'why?' (§§ 5, 15)

I want to suggest that (a) is part of an argument for (b); and understanding (a) and (b) together elucidate the special sense of 'why?' involved in (c). This special kind of explanation is then used to provide a non-causal account of the horizontal constitutive conditions of intentional action.

***(a) helps to support (b)*:** Compare two ways someone might explain *why the window broke*.

> *Causal Explanation:* The window broke because Homer threw the ball.
>
> *Intentional Explanation:* Homer intentionally broke the window by throwing the ball.

The causal story is intensionally transparent. Descriptions under which Homer's actions are intentional are intensionally opaque. Suppose that the window belongs to Ned, who unbeknownst to Homer has just left a box of pink frosted donuts on Homer's porch – one of a long series of neighbourly acts that Homer has failed to notice or left unacknowledged. It is true in just these circumstances that Homer broke the window, that Homer broke Ned's window and that Homer broke the window of his agreeable neighbour. But, in throwing the ball Homer might not have *intentionally* broken Ned's window; and in doing *this* he still might not have intentionally broken the window of his agreeable neighbour. There are circumstances in which the truth values of these statements about Homer's intentional action come apart from those of the first three statements.

Intuitively, these will have to do with Homer's psychology at the time he threw the ball – for example, with his relevant beliefs and desires. But if these must meet I-criteria, then they plausibly encompass more than any potentially causally efficient brain states or neural events occurring just at this time. Any content attributed to Homer makes sense of his action by fitting it into a network of interconnected psychological states and episodes over time and in a particular environmental context. For example, <Ned won't like *this*> thought of the act of throwing the ball, just before launching it at an impressive snowman Ned has built just in front of the window, is compatible with several overall interpretations, which will be more or less apt

depending on the how things are with Homer and how he is psychologically related to these temporally-extended circumstances. The interpretation isn't useful (and may not be available) in abstraction from this context, and so ascribing it (or some part of it) to a causally efficient neurophysiological mechanism likewise depends on circumstances that extend beyond the mechanism in both time and space. To interpret Homer as un-/intentionally breaking Ned's window, we must know something about "[t]he whole history of the incident" (Wittgenstein 1953, § 644, quoted in Anscombe 1957 § 25); and interpreting Homer as intentionally *breaking the window of his agreeable neighbour* will require us to know about (or lead us to expect) a different, probably longer and more surprising, history.

***(a) and (b) clarify (c)*:** What this brings out is the sense in which 'why?' asks a different question in relation to causal and intentional explanation. For Anscombe, to ask why Homer un/intentionally broke the window is not to ask for a psychological characterisation of the state or events *inside Homer's head* that led to his throwing the ball. Rather it is to request an interpretation, a characterisation of the state of mind that *Homer was in*, in light of which his actions are intelligible.[35]

Answers the question 'why?' asked in this sense give an agent's *reasons*, where "[r]oughly speaking, it establishes something as a reason if one argues against it" (Anscombe 1957, § 15). This means that answering 'why?' questions of this kind requires keeping track of history, context, motivation and belief all together. In a given situation, citing any one of these might serve to answer this kind of question. One might, for example, explain the ball's hitting the window as a mistake – Homer was aiming at the snowman. And we might take Homer's reasons for *trying to* hit the snowman as involving the belief that Ned has slighted him. In this case, pointing out Ned's long history of neighbourly behaviour might count as a reason against the action. Or it might only make things worse. If it did, this could count as evidence against the original interpretation and thereby reveal something else about Homer's psychology. But just what it would reveal is likely to depend on further features of the incident and its context. Crucially, this will involve a history of interactions between Ned and Homer that played out the way they did because of similar interpretive strategies that they used on each other.

[35] So, when interpretationists talk about psychological states, these are states that the subject is in, not states inside the agent: "on the interpretationist's view, propositional attitudes are not literally internal states. The concept of belief, say, is the concept of a property of a person: what is primary is the idea of a person believing that p; the idea of a kind of entity, S's belief that p, is derivative from, and dependent on, that." (Child 1994, p. 8). See also, Williamson (2009), for discussion Drayson (2018)

Interpretation is not only a useful tool for an external observer but is woven into the patterns of interaction that are the focus of interpretation.

The distinctive claim about horizontal constitution made by the philosophers of the personal level is that what accounts for the identity of a psychological phenomenon is a collection of circumstances relevant to its interpretation – it is the psychological phenomena it is because it stands where it does in the complex network of relationships tracked by psychological explanations. Subjects so situated will causally be affected, affecting and disposed to affect their environmental context in all kinds of ways that can be systematically (or at least non-arbitrarily) derived from the attributed psychological condition of the subject. For states with intentional content, this will involve the network of other contentful states attributable to the person and their normative relations to each other and to their environmental context. In part, this context is shaped by application of the same reasoning by other agents reciprocally interpreting each other. All of this might be characterised as a network of complex dispositions (though Anscombe points out that it needn't be); but since the behaviour in terms of which these dispositions are characterised (e.g., boasting, securing the admiration and envy of others) is described in intensional terms, this does not reduce psychological to causal relations specifiable in a more basic, physical vocabulary. Attributing content to belief, intention or perceptual states is a way of grasping the structure of that state's horizontal constitution conditions. This is why, for these philosophers, personal-level interpretation is *autonomous* – there is no way to hook onto this structure except from within the language of psychological explanation.

3) Subpersonal Explanation: What to Do when Your Spade is Turned

a) Indeterminacy and grounding

Reflecting on his earlier work, Dennett (2015, p. vi) attributes the distinction between personal and subpersonal levels of explanation to his "attempts to figure out what on earth Ryle was doing and how he could get away with it." We're now in a position to see what it was that Ryle and the other philosophers of the personal level had been, to Dennett's mind, getting away with. Without committing to an ontology of intrinsically psychological entities interacting in accordance with mechanical or paramechanical laws, they establish a framework for answering philosophical questions about intentionality and experience – about the thoughts, perceptions and actions of persons. According to this, rather than dealing with hypothesised inner entities, psychological explanation targets the patterns and structures manifested in interchange with an environmental and social context. This means that, in Schwitzgebel's (2002) terms, personal-level psychological explanations are 'superficial' rather than 'deep'; and *this* allows for a similarly

superficial way of treating constitutive conditions of psychological states. These are what make up the "real patterns" (Dennett 1991) in the psychological lives of subjects, acting and interacting with their environments. The approach proposed by these philosophers promises to transform intractable ontological questions about mental entities into questions of interpretation that could be asked and answered in a language closely related to that of ordinary psychological explanation.

But the cost of answering constitutive questions about mental states by appealing to features of (horizontal) psychological interpretation and explanation is that personal-level explanations come up against a limit. As we've seen, in order to be engaging in psychological interpretation at all, an interpreter must assume at least a minimal degree of rational holism and coherence. But, due to the open-endedness of psychological interpretation, appealing to personal-level explanatory norms will not constrain candidate interpretations to those that intuitively make *good* sense of the behaviour (Quine 1960; Gerrans 2004). Alongside ordinary explanations an inventive interpreter could always provide many more outlandish combinations of belief and desire that nonetheless ensure that behaviour is both environmentally appropriate and rational.[36] This gives us two puzzles:

> *Indeterminacy:* What makes the ordinary psychological interpretation a better explanation than any one of the outlandish-but-rational interpretations?
>
> *Grounding*: How can a complex disposition specified by an indefinitely large collection of psychological contents be grounded in a finite physical system?

On the reading I've been offering, the philosophers of the personal level answer the grounding problem by looking for constitutive conditions outside of the brain, in the practices of psychological interpretation and explanation (and, sometimes, their environmental context) that render behaviour intelligible as action. For an interpreter who is already a participant in these practices, this offers a solution to the problem of indeterminacy – this participation guides her to the interpretation that actually fits with the relevant conditions (Wittgenstein 1953, §198-202), but it also precludes a non-intensional statement of those conditions.

[36] For ease of exposition, the main text doesn't discuss the influence on this proposal of Quine's notion of 'radical translation'. Dennett frequently discusses Anscombe and Quine together (see, e.g. Dennett 1968, 1969, ch.9; 1978 ch. 1, 3; 1991), suggesting that although their responses to the puzzles raised by interpretation are radically different, Dennett draws similar lessons about psychological interpretation from both (for discussion, see Dennett 1978, 1991 Glock 1998).

The crux of Dennett's alternative proposal is that we can solve the indeterminacy puzzle by answering the grounding puzzle in neurocomputational terms. We can look to cognitive neuroscience "to provide a scientific explanation of the differences and similarities in what is the case in virtue of which different mental sentences [i.e. psychological interpretations of persons] are true or false" (Dennett 1969, p. 18) The idea is that rather than attempting to capture 'the whole history of the incident' in purely personal and interpersonal terms, interpreters can use the resources of psychological explanation to interpret the causal mechanisms that generate behaviour and mediate sensory input, identifying those that are relevant to the a given personal-level interpretation.

b) Subpersonal representation and neurocomputational function

How can cognitive scientists determine which neurophysiological events are relevant to psychological interpretation? Here's one answer. Start with the observation that the various ways in which persons are situated in relation to their environment and their own psychological histories (i.e., to their horizontal interpretive context) depend on their internal structural organisation, and in particular on the functional organisation of their brains. For any given psychological activity, some elements of this organisation will be causally responsible for the observable activity to which an interpretation is anchored. Call this kind of functionally individuated element a neural mechanism. The next idea is that components of neural mechanisms can be understood as carrying signals whose contents and roles can be understood by analogy to the properly psychological activities of *informing* about states of affairs or *commanding* that some action be performed. By ascribing contents to the messages carried by mechanism components in this way,[37] researchers make sense of internal functional organisation that would otherwise be intractably complex – reconstituting neural mechanisms as neurocomputational mechanisms.

Mechanism components, so understood, are individuated by truth, accuracy or satisfaction conditions – they realise and are charaterised in terms of representational content. Representational content, meanwhile, is in part the upshot of a relation between neurocomputational vehicles (patterns of neurophysiological activity apt for representational interpretation) and features of an environment such that the one acts as a signal for the other. For example, if inner neural events embody a structural correspondence with a feature of the environment, the resultant state's accuracy conditions are satisfied if the relevant

[37] Dennett (1969, p. 56) introduces the term 'neural vehicle' to refer to these components, in arguably the first use of the now-familiar sense of 'vehicle' in the literature. But see also Anscombe (1957, §59)

correspondence to the environmental feature obtains and fail to be when not (cf. Shea 2018, pp. 31-36). These informational contents are essential to, but insufficient for representational content for two linked reasons: relations of this kind may obtain without making any functional difference, and actually obtaining relations underdetermine any explanatorily useful content.[38] In Godfrey-Smith's (2006) phrase, what's relevant are the "exploitable relations" in which mechanism components stand to the environment. And exploitability is grounded in the causal influence that a mechanism has in the system.

Ultimately, talk of a representation's truth or falsity, success and failure conditions, tacitly relies on the persistence of the vehicle as a signal that can be used. Shea (2018) describes this as a "stabilisation of function", of which the main form is what Dennett (1969) already describes as a kind of 'neuronal Darwinism' (cf. Edelman 1987):

> [The] link between each bit of Intentional interpretation and its extensional foundation [i.e. vehicle] is a hypothesis, or series of hypotheses describing the [neuro]evolutionary source of the fortuitously propitious arrangement in virtue of which the systems operation in this instance makes sense. These hypotheses are required to make sense of the appropriateness which is presupposed by the Intentional interpretation, but which requires a genealogy from the standpoint of the extensional, physical [i.e. neurophysiological] theory. (Dennett 1969 p. 80).

Success and accuracy conditions are normative for representational mechanisms because failure to reflect them would, outside of certain bounds, cause a mechanisms component to fail to represent anything. If a vehicle fails to represent, it is not reproduced in the system. On the other hand, iterative reinforced success explains representational success.

Connecting this methodology to the work of the philosophers of the personal level, Dennett points out that this is an application to *subpersonal* mechanism components of a psychological explanatory strategy grounded in the interpretation of whole agents. At least at first, it's only by taking it that the contents of messages internal to the brain's processing are subject to norms of consistency and environmental appropriateness that they can be understood as messages or

[38] Note that, we can't appeal to the strongest correlation to address this underdetermination problem (e.g., to fix singular contents for representations of objects and their properties), since proximal events will always be more strongly correlated – if correlation strength fixed content then no vehicle could represent anything in the environment. Indeed, it's a defining feature of *perceptual* representational states that they signal their distal target and *not* these causal intermediaries (Millikan 2001; Shea 2018).

commands.[39] At the sensory or motor efferent peripheries, interpreting a neurophysiological event in this way (e.g., as the command 'contract now, muscle!'; Dennett 1969, p. 79) gives no greater explanatory purchase on its functional contribution to the production of behaviour (e.g., causing a muscle to contract) than seeing it as manifesting a purely causal disposition. More centrally, however,

> characterizing a state or event or structure not only as a physical entity operating under certain causal conditions but also as, for example, a specification of a goal or description of the environment or order to perform a certain task would be virtually the only way of 'making sense' of neural organisation (p. 80)

While norms support indispensable explanatory heuristics at the personal level, these lose their salience as subpersonal explanation is applied to lower levels of functional organisation in the brain. Just as the 'command' a motorneuron gives to a muscle fibre doesn't need to be interpreted against the background of beliefs and desires attributed to this subcomponent of the overall mechanism, the functional contribution of any fundamental component can, in principle, always be understood causally. Taking the initial interpretation for granted at first, the subpersonal questions ask what underlies the patterns in persons' behaviour that manifest just those attributed contentful states – highlighting some and disregarding other parts of the network of neurophysiological causes of what they say and do. Followed through to the lowest subpersonal level, a non-intensional statement of the constitutive conditions of the most basic kinds of subpersonal psychological phenomena resolves the grounding problem by dissolving attributed content. At this level, there is no problem of trying to fit indefinitely many psychological contents into a physical system, because none of the attributed contents are in any of the sub-systems at the lowest levels of explanation – they are in their context, which, by now includes the neurocomputational mechanisms that any one sub-system interacts with. At this point, content attribution really is heuristic in McDowell's sense: to the extent it is useful, it is only in that it offers a shorthand capturing the mechanism component's relations to wider processes. Content attribution captures the 'real patterns' in the relations between a person's brain, body and world that fit her into an interpersonal, psychological context. Though psychological terms are required to capture the patterns, they are grounded in nothing more – or less – mysterious than the physical structure of the world. So, while the personal-level

[39] Note that, at this stage in the argument, Dennett is making a purely methodological point – understanding the contribution of subpersonal mechanisms to overall psychological functions relies on interpretation. This is prior to and independent of the interpretationist claim that these practices provide identity conditions.

conditions of interpretation don't reduce to 'personal-level' event causation, neither do they introduce a free-floating domain of autonomous non-causal explanation.

c) Model-based fMRI and subpersonal explanation

Model-based fMRI studies illustrate one way in which Dennett's basic picture of subpersonal explanation is borne out in contemporary cognitive science. The key assumption behind all fMRI studies is that neural activity correlates with blood flow to, and so the blood-oxygen-level-dependent (BOLD) signal emitted by, relatively small areas of the brain due to the greater metabolic demand generated by this activity (Hilman 2014). Model-based approaches also assume that differences in neural activity correspond to differences in values computed by algorithmic models of cognitive functions derived from behavioural data (Gläscher & O'Doherty 2010). This means that model-based studies involve two distinct steps. The first uses behavioural data to establish a set of candidate algorithms that could be producing the observed behaviour; the second selects from candidate algorithms from the first stage, fitting these to observed patterns in the BOLD signal. The idea is that, averaging across recordings from multiple trials and multiple participants, the algorithm that best models the neural mechanisms responsible for the behaviour will fit the BOLD signal data better than other candidate algorithms.

It's important to note that the 'observed patterns of behaviour' mentioned in the first step are not patterns of uninterpreted bodily movement. The only algorithms that are available for consideration are those that make sense of what the experimental participants are doing. The relevant behavioural data consist of actions interpreted as falling into different categories, and these categories are determined by the context that the actions are performed in. While these are realised by movements which exhibit systematic differences at the level of gross-bodily (i.e., uninterpreted) motion, these differences do not play a systematic role in explanation except in as far as they are situated within a context assumed (and so not explicitly specified) by the researchers.

Psychological interpretation plays a central and ineliminable role in this kind of cognitive scientific explanation. Though personal-level phenomena do not feature as model components, the algorithmic models proposed at the first explanatory stage are useful only as models of activities identified by reference to their position in a network of relations (to other psychological phenomena, to the environment) that the early interpretationists appeal to in horizontal constitutive explanations. By using neurophysiological data (the BOLD signal) and anatomical considerations about neural architecture to guide model selection, researchers map these models onto neural mechanisms by assigning contents specified by the model of

mechanism components. As they are mapped onto neural structure, these interpret underlying causal interactions as realising distinct subfunctions. These subfunctions, in turn, constitute the overall psychological capacity associated with a personal level state or states.

With this outline established, the following section will make it more precise in two ways. The first of these will bring out more of the general features of subpersonal explanation as it was proposed by Dennett (1969) and refined, with the development of cognitive science and its philosophy, over the next two decades. The second introduces a constraint on acceptable explanatory mappings. Adopting this constraint is what gives us the isomorphism principle.

4) Neurocomputational Mechanisms and Isomorphic Explanation

a) Task analysis and functionalism

The first step of the model-based studies described above is an example of a more general methodology in cognitive science called 'task analysis' (Chipman et al. 2000) or 'functional analysis' (Cummins 2000). This has its roots in the 'cognitive revolution' (Miller 2003) in psychology, and was elaborated by cognitive scientists and philosophers, including Cummins (1975), Fodor (1968, 1975), Lycan (1981, 1991) and Dennett himself (1978, 1987) through the 1970s and '80s. Task analysis characterises a psychological capacity as a set of component operations, or subfunctions. This approach puts algorithms at the centre of psychological explanation. Generalised, it entails that "the paradigmatic psychological theory is a list of instructions for producing behavior" (Fodor 1968, p. 630). And, as we've already seen, producing the right kind of behaviour is how you get to be interpretable as being in a psychological state.

Here's a toy example. Take a psychological capacity, Ψ, of a subject, S. Assume that Ψ is characterised by S's relations to the environment, so that when S exercises the capacity S is interpretable as being in a corresponding psychological state. A task analysis of Ψ might involve three subfunctions (ϕ_1, ϕ_2, ϕ_3) that jointly realise the capacity:

Function, Ψ	
is realised by subfunctions, ϕ	ϕ_1, ϕ_2, ϕ_3

Table 2.1

Since ordinary personal-level interpretation of S as being in a state or performing an activity associated with Ψ (e.g., 'S is Ψ-ing') does not always distinguish between the ways this is realised, we might wonder if by saying that S is Ψ-ing we are also saying that S is, for example, ϕ_1-ing? Depending on what Ψ is, we might be drawn to different answers. Solving an arithmetic

problem, 7x12=?, could involve either calculating or remembering a multiplication table. Intuitively, what makes it the case that S is either *calculating* or *remembering* the answer depends on which algorithm is actually realised in her brain, but up to a certain level of detail we are also happy to attribute the stages of that algorithm to S herself (e.g., first multiplying 10 by 7, etc.). What happens after we reach this level? A task analysis for remembering a multiplication table is likely to involve intelligent activities (e.g., *retrieving* and *interpreting* stored information) that we wouldn't attribute to S in the same way.

Compare two questions. Q_1*:* 'how did you know the answer?' (*A:* 'I remembered my times-tables') with Q_2*:* 'how did you do that?'. This is where the philosophers of the personal level say that our spade is turned – an answer to Q_2 can either be 'I just can' or it can name non-psychological causes of behaviour (Wittgenstein 1953, § 217). Task analysis proposes a way to keep on digging by analytically decomposing the activity. If I-criteria are determined by subpersonal facts about neural mechanisms, this more detailed interpretation can be constrained by these facts (e.g., in the way illustrated by model-based fMRI research above).

Some early accounts of how this next level of psychological explanation is supposed to work rely on the notion of an explanatory 'homunculus' – a metaphorical inner agent who carries out psychological functions (see, e.g., Lycan 1981, 1991). Worries about homuncular explanation in psychology had motivated both psychological behaviourists, like Skinner (1957), and the philosophers of the personal level. As they saw it, a giving a homuncular answer to Q_2 would be to say something like 'my memory *remembered* the solution', which either defers an answer of one of the first two kinds or starts off a regress. But if tasks can be decomposed, why not homunculi? The proponents of 'homuncular functionalism' argued that cognitive scientists should

> account for the subject's intelligent activity, not by idly positing a single homunculus within that subject whose job it simply is to perform that activity, but by reference to a collaborative team of homunculi, whose members are individually more specialized and less talented (Lycan 1991, p. 259)

So, subfunctions could be assigned to subpersonal homunculi (sometimes, 'subagencies'). The realisation of a psychological capacity could then be explained as the coordinated performance by the members of this collaborative team of their individual tasks. Now we have two levels of psychological interpretation, one pertaining to the person, S, and the other to subpersonal subagencies, s:

Personal, Ψ	S Ψ-ing
Subpersonal, ϕ	s_1 ϕ_1-ing, s_2 ϕ_2-ing, s_3 ϕ_3-ing

Table 2.2

While Dennett (1969) had warned that cognitive scientists should not "replace the little man in the head with a committee", by (1978) he too had embraced an approach on which researchers start

> with a specification of a whole person or cognitive organism […] and then [decompose] that intentional system into an organisation of subsystems, each of which could itself be viewed as an intentional system (with its own specialised beliefs and desires) and hence formally a homunculus (p. 134)

Dennett emphasises the same two features mentioned by Lycan (1991) as relevant to this lower level of psychological interpretation: subpersonal homunculi are specialised (they have distinctive roles, not performed by the whole person, in light of which their functional contributions to the overall system are intelligible), and they are less intelligent than whole persons. This means that successive rounds of task analysis eventually 'dissolve' or 'discharge' these homunculi by specifying tasks at a low enough level that their possible contribution could be understand as the upshot of mechanistic causal-explanatory principles.

b) Mechanistic explanation and the isomorphism constraint

So far, it's not clear how subpersonal agencies, introduced via the notion of an explanatory homunculus, are related to the methodology described in (§2.3b). This contains no reference to homunculi, or to subpersonal interpretation. One natural way to relate the two is to associate homunculi with mechanisms: task analysis provides an abstract model of the psychological capacities that cognitive neuroscientists map to neurophysiological mechanisms.

At this point, it will be helpful to say more about what a mechanism is. Throughout the above discussion, I've been assuming a now-standard definition,[40] according to which a mechanism is

[40] The definition is most strongly associated with the 'new mechanism' in philosophy of science (see also, e.g., Machamer et al., 2000; Bechtel 2008; Craver 2007). For these theorists, mechanistic explanation does not necessarily involve appeal to deterministic laws or to discrete 'mechanical' structures, along the lines of 'old' mechanistic accounts of natural phenomena that emerged from European early modern science (Craver & Tabery 2015) and which arguably still make up the background to Dennett's earlier proposals.

> a structure performing a function in virtue of its component parts, component operations, and their organisation. The orchestrated functioning of the mechanism is responsible for one or more phenomena (Bechtel & Abrahamsen 2005, p. 421)

A mechanistic explanation, then, involves identifying parts of mechanisms, their relations and their functions. Mechanistic explanation, in this sense, is ubiquitous in cognitive science (ibid.; Craver & Bechtel 2008), and more generally it is a feature of many of the 'special sciences' that deal with phenomena at different explanatory levels (Fodor 1975; Bechtel & Abrahamsen 2005). Mechanistic explanation involves associating components and functions with each other, so that researchers can understand a phenomenon by understanding the orchestrated functioning of components of its underlying mechanism. For example, interventions on mechanism components will have effects systematically predictable from an appreciation of the functions of those components. So, the relevant notion of 'responsibility' is a kind of vertical explanatory dependence: mechanisms constitute phenomena and thereby causally explain effects brought about by the coordination of its components. Crucially: although causal interactions compose the mechanism, all causal relations are intralevel. A component's activity might have a systematic effect on higher level phenomena – that it sometimes does is part of what is involved in its being a mechanistic component – but this effect is necessarily mediated by its relations to other components and extra-mechanism elements. The patterns of interaction of elements at the lower level non-causally determine features of the phenomenon it mechanistically explains either synchronically (mechanistic constitution) or diachronically, via mechanistically mediated effects.

We can now elaborate on the basic picture illustrated in tables (1, 2) to introduce the isomorphism constraint. First add some more detail to the idealised picture of functional analysis from table (2). Introduce an element, x, that is not part of the task analysis for Ψ given in table (1), but which causally mediates the behavioural output that characterises Ψ. Assume that each instance of ϕ-ing can be characterised by two relations, R, that uniquely hold between the s that is ϕ-ing and other elements. Any s can be in a unique R to any other s or to any x. Take the symbol '→' to indicate which relations are involved in any subpersonal mechanism description. We can use this expanded vocabulary to give the following subpersonal task analysis for S Ψ-ing:

Personal, Ψ	S Ψ-ing
Subpersonal, ϕ	s_1 ϕ_1-ing → R s_2, R s_3 s_2 ϕ_2-ing → R s_3, R x_4 s_3 ϕ_3-ing → R s_2, R x_4

Table 2.3

This turns the unstructured list from table (2) into something more like a structured set that could be mapped onto other models. For example, elements of the task analysis could be mapped one to one onto components of a model showing how causal mechanisms in the brain interact to produce behaviour. Doing so would ensure explanations based in these models meet what Hurley describes as an 'isomorphism requirement' (Hurley 1998, ch. 10; see also Hurley 2003a) on explanatory mappings between levels. An isomorphic mapping is one for which mapping functions can be inverted. On this example, it is one on which the inverse of a function that maps a subpersonal agency onto a mechanism component uniquely maps that component onto the same agency. If we identify each of these subagencies in terms of its relations, then these will have to be preserved in the causal structure of the isomorphic explanatory mechanism. The diagram below (figure 1) adapts Craver's (2007) general model of explanatory mechanisms to give a picture of this kind of mapping.

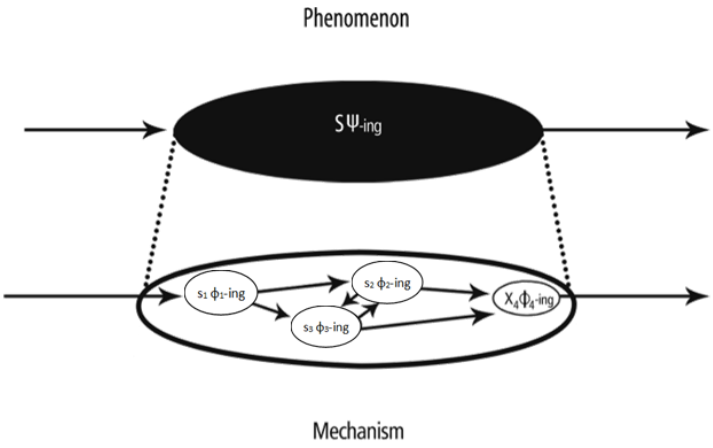

Figure 2.1

This isomorphically maps the task analysis given in table (3) to the depiction of a causal mechanism. The arrows between components of the mechanism (the smaller ovals) represent causal relations. Mapping the task analysis from table (3) onto these interprets them as realising instances of the functional relation, R. If we ignore extra-mechanism relations (represented by arrows passing out of the larger ovals) it's not hard to see that one could formally characterise a mapping function from elements of a more realistic task analysis to mechanism components.

When this kind of relation between models holds, we have an interlevel explanatory isomorphism.

Note that, so far, talk of 'isomorphisms' refers to relations between models, tables and lists. That is, the relation of isomorphism is a relation between semantic items used in cognitive scientific explanations of psychological phenomena and not the phenomena themselves. I take this to be the primary meaning of 'explanatory isomorphism' in Hurley (1998), and it is the meaning I adopt. Explanatory isomorphisms are thus between elements in models that can be given a mathematical specification – where the specification of counterpart items in two models match in a way that allows for two-way mappings between the models, these models are isomorphic. When models are isomorphic and this isomorphism is put to some kind of explanatory use, they have an explanatory isomorphism. Of course, models in cognitive science are used to explain real world phenomena like neurological and psychological mechanisms. So, the explanatory benefit that isomorphism is meant to grant carries over, ideally, to phenomena modelled: explanatory isomorphisms between models are meant to support explanations of these phenomena. This licences a secondary use in Hurley's work of isomorphism as a way of describing relations between *phenomena* at 'personal' and 'subpersonal' levels whose features and dynamics are adequately captured in models that exhibit the kind of isomorphism set out above. This is harmless and does not adopt any commitments beyond those already taken up just as long as it is understood as an extension of the primary sense that applies to models. Introducing isomorphisms as a distinct kind of relation between phenomena, however, begs important questions – it seems to presuppose the distinct 'levels of reality' complained of by Rupert (2018) – and inappropriately applies the mathematical notion of isomorphism to non-mathematical objects. So, although I follow Hurley in talking of isomorphisms between phenomena, it is important to bear in mind that this is a secondary use of the term. It is intended in all cases to be justified in virtue of the appropriateness of explanations involving isomorphic models in the primary sense.

Importantly, not every element of an isomorphic model in cognitive science needs to be attributed content (i.e., interpreted in as behaving in a content-guided way). Returning to figure (2.1): the rightmost component, x_4, can be understood as producing its output 'mechanically' given the input it receives from s_2 and s_3 and not representing any content specified by the algorithm. For example, one might think of this as a mechanism component responsible for a low-level 'contract now!' motor command. Both x_4 and the other, subpersonal, mechanism components are subject to further mechanistic decomposition, which allows their functionally interpreted causal dispositions to be explained in terms of the interaction of components at a

lower level of explanation. These might be content-ascriptive in the way just outlined; but at some point in this process, they won't be. This captures how, for Dennett, homuncular explanations involving intentional content attributed are discharged into simple event-causal explanations.

5) Isomorphic and Non-Isomorphic Styles of Subpersonal Explanation

a) Case study: explanatory isomorphisms and visual migraine aura

To help make the previous section's discussion a bit more concrete, it will be useful to consider a case in which explanatory isomorphisms shed light on a real psychological phenomenon. About one third of people with migraine (or 'migraineurs') experience some form of sensory disturbance, starting before the onset of headache and other symptoms. A recent review of the neuroscientific and medical literature on these 'migraine auras' characterises them as "the remarkable complex, symptom-rich, pleomorphic neurological phenomenon present in some [migraine] attacks" (Hadjikhani & Vincent. 2021). Visual migraine aura and the development of scientific accounts of its relation to underlying neurophysiological processes make a useful case-study, illustrating how isomorphic subpersonal explanation might work and why explanatory isomorphisms can be so appealing. Since visual migraine auras are fairly common, they also offer a helpful supplement to philosophical discussion that is often driven by much more unusual conditions, such as 'blindsight' (Kentridge et al. 1999, 2008) and phenomena of 'hemispatial neglect' (Parton et al. 2004).

A common form of aura presents as a 'scintillating scotoma', a slowly expanding patch of brightly flickering visual space that obstructs normal vision. We owe this term to a tradition of first-person reporting by 18th and 19th century scientists (Eady 2009; Foxhall 2019), who documented and sometimes meticulously illustrated their aura hallucinations. One of the earliest accounts of these distinctive visual symptoms describes them as "a singular kind of glimmering in the sight" in which "objects change their apparent position surrounded by luminous angles, like those of a fortification." (Fothergill 1778, quoted in Foxhall 2019, p. 113). These features, which will be familiar to many people who experience migraine (as well as some people who don't[41]), are reproduced by illustrations like those that accompany Hubert Airy's (1870) classic discussion of aura.[42] Classic scintillating scotoma expands gradually across a lateralised patch of the visual

[41] While some people report a similar phenomenology, without accompanying headache, I'll follow the mainstream of the contemporary literature in referring to these together as *migraine* auras (Hadjikhani & Vincent. 2021; O'Hare et al. 2021).

[42] Foxhall describes these as "arguably some of the most beautiful imagery in the history of medicine" (Foxhall 2014, p. 364). They really are quite beautiful, and you can see them, as well as several other

field and is, as in Fothergill's and Airy's descriptions, often accompanied by angular visual hallucinations ('fortification spectra') that form a crescent on the expanding edge of the scotoma. This characteristic phenomenology fascinated early neurologists and contributed to an air of mystique around the condition, believed to be caused by the excess of 'nervous strain' displayed in these symptoms.[43] Auras, the thought was, afforded migraineurs "a veritable "Photograph" of a morbid process going on in the brain" (Airy 1870, p. 247).

Contemporary research into migraine aura identifies spreading cortical depression (SCD) as the physiological mechanism that most likely underlies these symptoms. Strikingly, the explanation of aura as a manifestation of SCD seem to bear out this assessment – directly linking this aspect of aura phenomenology with precisely localised disruptions in the neurocomputational mechanisms of vision. SCD is a physiological process that result in characteristic pattern of disruption in the ordinary function of affected neural areas. As one review notes, this disruption takes the form of a wave, gradually spreading across the cortical surface:

> At the core of [SCD] is a rapid and nearly complete depolarization of a sizable population of brain cells with massive redistribution of ions between intracellular and extracellular compartments, which evolves as a regenerative, "all-or-none" type process and propagates in the manner of a wave through gray matter. (Goadsby et al. 2002, p. 260)

SCD itself was first described by Leão (1944) and, except in cases of stroke or traumatic brain injury (Strong et al. 2002) has not been directly observed in humans. Physiological evidence for SCD during migraine aura and in other context is derived from functional imaging modalities, like fMRI and EEG (Hadjikhani et al. 2001; Hadjikhani & Vincent 2019). So, as with many cognitive constructs, what we are presented with is indirect evidence for a mechanistic interpretation of

---

19th Century depictions of visual aura, here: https://publicdomainreview.org/collection/visualizing-migraines

[43] As Foxhall's (2014, 2019) work on the medical and cultural history of migraine narrates, the perceived centrality of unusual visual and cognitive phenomenology to a condition whose symptoms had prior to the 19th century been characterised largely in terms of its associated patterns of pain and bodily discomfort (as 'hemicrania' [lit. half-head], from which the term 'migraine' is derived; or the 'bilious-' or 'sick-headache') owes as much to the fascination these symptoms held for early neurologists as it does to their clinical significance or prevalence. This fascination had a gendered dimension: the nervous strain displayed in migraine aura was frequently associated with 'male genius', and so migraine was characterised as a typically masculine disorder. In fact, modern prevalence studies report that women migraineurs outnumber men by a ratio of about 3:1, with *more* women overall than men reporting auras (though as a slightly lower proportion). So, it might be worth noting that for a majority of all people who experience migraine, attacks are rarely, or not at all, accompanied by aura symptoms (Goadsby 2017; O'Hare et al. 2021).

a complex dynamic phenomenon rather than a concrete model with independently specified components.[44]

What is this evidence and how does SCD explain aura symptoms? Neurons in early visual areas of human and many other animal visual systems are spatially arranged across the cortex in columns, with nearby neurons responsible for processing input from the same or nearby retinal cells (by way of intermediary subcortical processing stages – in mammals, most importantly LGN – and with significant feedback input from cortical areas). The overall effect of this arrangement is that areas of cortex can be systematically mapped onto retinal receptive fields (retinotopic mapping), and so onto areas of the visual field. When these are affected by neurophysiological phenomena like SCD, reported subjective effects can be systematically related to the affected areas. When SCD is focalised on the early visual cortex, the slow spread of depolarisation is accompanied by the growing 'scintillating' patch in the visual field. Since the process typically unfolds over about twenty minutes, it falls within the temporal resolution of fMRI. The dynamics of SCD seem to track reported phenomenology in real time, with the affected area of the retinotopic map corresponding to the area of visual space affected by the aura symptoms (Hadjikhani et al. 2001; O'Hare et al. 2021). Moreover, the character of reported phenomenology tracks the selective processing roles standardly assigned to visual areas.

To reinforce the point, let's consider a particular instance of isomorphic explanation in more detail. Here, I'm going to focus on how these are related to the underlying mechanisms of CSD in a single patient study (Hansen et al. 2013). Because of the unusual methodology of this study, it gives a particularly clear example of isomorphic subpersonal explanation (see also Wilkinson 2021). However, since the study doesn't involve functional data, it can only be understood as localising the dynamic neurophysiological mechanisms of migraine aura in the context of other work with fMRI (e.g., Hadjikhani et al. 2001)

As Hansen and colleagues report, their experimental participant, PV

> made consistent and detailed drawings of his migraine visual auras over nearly two decades and created a database of 1000 aura attacks. This database includes detailed maps of his visual perception of hundreds of his migraine auras, noting the precise site of initiation and pattern of spread throughout his visual fields. (p. 3590)

[44] Given the outsized impact that evidence from imaging modalities has on consumers of neuro- and cognitive scientific literatures (McCabe & Castel 2008), this bears emphasising when presenting this kind of evidence.

One of these images is reproduced in (figure 2.2, A). The authors' description of their data collection places this study squarely into the tradition of begun by Airey (1870) and his 19th Century colleagues. Interpreted against the background of previous findings about the spatial arrangement and functional specificity of visual cortex, PV's regimen of structured introspection also yields material for isomorphic subpersonal explanation. After the corpus of images was digitised and analysed, it was mapped onto structural data from MRI and interpreted in terms of the SCD theory. For example, the researchers use their analysis of their dataset to reach conclusions about the temporal and spatial dynamics of SCD events associated with reported phenomenology:

> Based on the assumption that the underlying cortical wave travelled in a direct path along the calcarine sulcus and crossed the entire primary visual cortex, from the occipital pole to the most rostral part of the striate cortex at the confluence of the calcarine and parieto-occipital sulcus (in this individual 62 mm), the speed of the underlying cortical event was estimated to range from 2–3 mm/min (Hansen et al. 2013, p. 3591)

More important than the details of how these results are generated, is the style of explanation offered. What should strike us is that precisely quantified models from personal-level behavioural data are mapped on to subpersonal models (informed by anatomical data and prior theory). Explanations in this style move smoothly between subpersonal and personal levels because of the systematic way in which models at each level (respectively, PV's recordings of his migraine aura attacks and standard models of early visual cortex organisation) correspond to each other. A further example is provided by figure (2.2) (for a detailed description, see *this note*[45]):

[45] **(A)** reproduces a representative sketch by PV of the progression of a typical migraine aura. The numbers correspond to the time into the attack by minute. (1-7) illustrate the transit of a positive symptom through visual space until it disappears at (7); (12-18) correspond to scotoma, which gives way to another wave of positive symptoms starting at (17). **(B)** shows a schematic representation of the retinotopic map of human V1 for right visual field/left visual cortex. The grey shaded area is the 'monocular crescent' (the area corresponding to visual input from the contralateral eye only) the black dot is the blind spot. **(C)** shows a hypothesised path across the visual cortex of SCD associated with symptoms recorded in **(A)**.

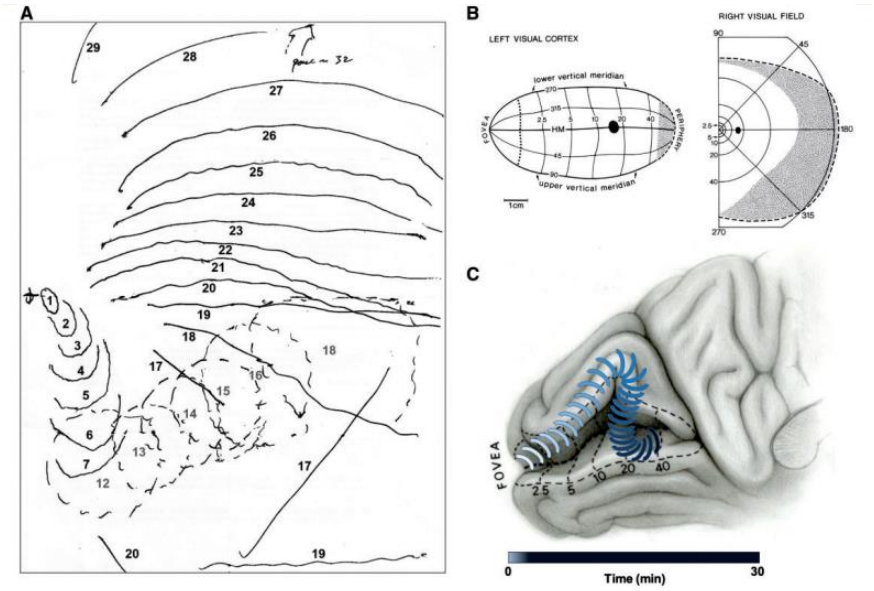

Figure 2.2

The isomorphic character of the explanation is particularly clear here because it's depicted visually. Representations like figure (2.2) support structure preserving mappings onto abstract models of the functional architecture of the brain and onto anatomical (Hansen et al. 2013) and physiological (Hadjikhani et al. 2001) data. These explanatory mappings allow for details of a personal level perceptual phenomenon (i.e., the migraine aura) to be characterised in detail by specifying their underlying neurocomputational and neurophysiological mechanisms.

This framework involves mechanistic explanations that are not subpersonal explanations: at the cellular level, having to do with ion-channels involved in polarisation; at the neuroanatomical level, having to do with the propagation of the SCD wave across the visual cortex. But since the relevant structural organisation is interpreted as carrying signals from the retina that part-constitute functionally analysable states of the person (PV's visual percepts and distortions of these) it is also subpersonal explanation. That these levels of mechanistic organisation are so smoothly integrated in this case is a sign of the progress of visual neuroscience in characterising more fundamental mechanisms (Trenholm & Krishnaswamy 2020). Nonetheless, in as far as the early visual system is understood as contributing to the personal-level capacity to enjoy visual experience (or to suffer disturbances in the exercise of this capacity), this does not speak against the explanatory usefulness of content-ascriptive functional decomposition.

How does this help us answer the constitutive question: *what is it to experience a visual migraine aura?* The availability of isomorphic explanations of migraine aura, which the above discussion

points towards, suggests that the horizontal and vertical constitution conditions of at least some psychological phenomena come together. Identifying mechanisms and tracking their interactions at subpersonal and neurophysiological levels enables precise prediction and explanation of a corresponding feature of a paradigmatically personal-level phenomenon, an introspectable and reportable visual experience. This is where I want to suggest that Airey's (1870) description the experience of aura as a "veritable photograph" of inner neural processes gets something importantly right. Visual aura makes the disruption of neurocomputational capacities perceptually salient to its subject, and what makes one interpretable as having an experience of this kind (against the background of otherwise normal criteria for interpretability) can be wholly specifiable in terms of the underlying neural mechanisms. In this case, migraine researchers may have, in Dennett's (1969) phrase, "pinpointed the conditions" that rendered PV's precise descriptions of his experiences accurate.

b) Interpretationism and the two-level interdependence view

We can now introduce the two-level interdependence view. This gives us a way of accounting for the constitutive relevance of mechanisms (and disruptions in their ordinary functioning) that this case study illustrates. Aura appearances are constituted by the disruptive effects in neurocomputational processes that SCD brings about. Attention to the structure of their interlevel explanation, we can see why. When we account for their phenomenal and features at the personal level, we do not have to appeal to causal processes external to the underlying neurocomputational mechanisms. In the next chapter, I'll rely on the same account to argue that ordinary, perspectival, object perception is not like this – object appearances are not constituted by neurocomputational mechanisms, although these processes still play a constitutive role in the structure of experience. We begin to understand the difference when we understand the distinction between the isomorphic subpersonal explanations appropriate to the first case and the non-isomorphic explanations appropriate to the second.

Hurley argues that we should reject not isomorphic explanations but rather the assumption that explanation must be isomorphic. What's more, she presents the commitment to non-isomorphism as a way of avoiding an opposing extreme, the denial of interlevel explanatory relevance:

> Clarity and progress [in cognitive science] are served by distinguishing levels and framing issues explicitly at a given level, or as concerning interlevel relations. Sliding between levels on a priori assumptions of isomorphism is unjustified. Nevertheless, one level can shed light on another. (Hurley 2008, p. 20)

So, for a model at the subpersonal level of content-individuated neurocomputational mechanisms:

> [w]e can look "down" a level, seeking neural implementations of aspects of [a modelled phenomenon's] functional architecture, or "up" a level, considering what [that] functional architecture would enable persons to do. (ibid.)

Although explanatory isomorphisms shouldn't be assumed a priori, they might be found to hold in some cases. Nonetheless, Hurley has little to say about what these cases are. Her discussion of isomorphism is mostly critical and negative. Explanatory isomorphisms are what fail to hold in the case of perception and action – they are what give rise to confusion in attempts to give interlevel constitutive explanations of these capacities. We might think this is unsurprising. Her focus is on perception and action, and it's a consequence of her view that finding isomorphisms elsewhere will be a matter of empirical investigation. Nonetheless, this has the unfortunate effect of obscuring the principles behind the distinction between isomorphic and non-isomorphic styles of explanation. It can't simply be that content attribution calls on wider context, because in this case we could *rule out* isomorphism in subpersonal explanations in advance.

To reconstruct a more balanced view of the place of explanatory isomorphisms in Hurley's overall picture, we should nonetheless start with these negative comments. When critiquing explanatory isomorphisms of perception and action, Hurley talks about relations between personal-level and subpersonal capacities (perception and action on one hand, motor and sensory information processing on the other), boundaries (of the subject and agent/ of neurocomputational mechanisms) and properties (of contents/ of vehicles). I've suggested above that we gain a unifying perspective when we think of these as elements of explanatory models pitched at different levels of mechanistic explanation. Mechanisms realise capacities, and neurocomputational mechanisms do this by processing information in a way that justifies content attribution at multiple mechanistic levels. Boundaries of mechanisms, meanwhile can be determined by input-output relations relevant to determining their functions. Explanatory isomorphisms are interlevel relations where identical contents, boundaries, or causal structure are found across levels. So understood, the migraine aura case involves interlevel isomorphism of all these kinds.

To see why, let's start at the personal level by considering (part of) what goes wrong with vision during migraine aura *within the perspective* of someone experiencing aura. The relevant notion of perspective is semi-technical, and it will come to play an important role in the perspectival

acquaintance account developed in later chapters. So, it will be useful to provide an brief characterisation at this point in the discussion. 'Perspective' here refers to a set of relations holding between things experienced by the subject (and so contents ascribable to that subject). It needn't suggest something that can only be understood 'from the inside', or some feature of experience that is ineffable, incorrigibly accessible, or essentially private. Conversely, denying that perceptual experience has these features doesn't amount to denying that it has perspectival structure. Perspective can be characterised in terms of the subject's perceptual capacities in a way that is neutral about these claims. Hurley (1998, p. 86) gives a characterisation of this kind:

> [H]aving a perspective means that what you experience and perceive depends systematically on what you do, as well as vice versa. Moreover, it involves your keeping track […] of the interdependence between what is perceived and what is done, and hence awareness of your own agency. (Hurley 1998, p. 86)

Having a perspective, in this sense, is bound up with being interpretable as a subject and an agent. A claim that Hurley repeatedly comes back to in setting out her positive account is that perception and agency are interdependent at the personal level. Here then, "awareness of your own agency" stands not only for a subjective capacity for õself-consciousness, but also for what that subjective awareness is an awareness *of* – namely, of perspectival structure in the experiences of a perceiving and acting subject. This structure, crucial to such an agent's self-interpretation as the kind of creature it is, is also the target of third-personal interpretation that might specify the conditions for subjecthood and agency. In order to understand a subject as perceiving the world, one must be able to think of her as thereby able to act, or intend to act, on the world she perceives.

This fact about what is involved in interpreting perceivers and agents as such, parallels the claims about psychology made by philosophers of the personal level. A consequence is that we can't attribute content to a subject's experience without understanding the range of intentions that she could form in light of this content. This is why agency is important. While we might think of perspective as a purely spatial notion – as a point from which we could imagine projecting geometric perspective lines (see Gregory 2013, ch. 2) – this is not the notion called on by interpersonal interpretation. For example, an agent who exhibited non-standard but systematic movements in response to stimuli might be interpretable as having a distorted perspective on the environment she is in (for example, as misperceiving the real environment as an 'El Greco world'; see Hurley 1998, ch. 8). In such cases, we would rely not on the features of the objective

spatial perspective' that someone in her position would occupy (or not only on these), but on her interactions with the objective world to make sense of her subjective perspective on it. While such distorted perspectives can be constructed 'objectively', attributing them to subjects requires that we take perception-action interdependencies into account.

It's thus a condition on being interpretable as an agent that there are *some* systematic relations between perception and intention, and so agency. An agent can only interpret herself as such if she keeps track of these relations in her own case. Likewise, an external interpreter must also grasp the differences that changes in an object or an agent's bodily orientation would make given that agents (changeable) internal states. In the limit case, we lose our grip on what it means to say that someone is perceptually aware of an object in front of her if we deny that perceiving that object could make a difference to the intentions she could form.[46]

The foregoing sketches an ideal model of the relationship between perception and action at the personal level of explanation.[47] Let's call this *perspectival interpretationism*:

> *Perspectival interpretationism:* Perception and action are co-constitutive. It is essential to being a perceiver that one is an agent – and *vice versa* – and being a perceiver-agent means having a perspective. Since perspectival structure is tracked by psychological interpretation, subjects and agents are subject to I-criteria.

With this sketch in place, we can characterise the effect that migraine aura has on the perception-action dependencies that constitute the perspective of someone experiencing aura. The evidence reviewed in the previous section (§2.5a) tells us that the area of visual space affected by aura symptoms directly reflects the area of cortex affected by SCD. Since the systematic dependencies that Hurley describes typically run through the environment, they do not show up in the affected area. This is because the channel for information about the environment passing through the corresponding area of early visual cortex has been partly or

[46] The converse is not obvious and possibly not true. It seems at least plausible that we can make sense of the idea of an object affecting the range of intentions available to a subject without her being perceptually aware of that object. Some discussion of blindsight understand it as this kind of case (for discussion, see Marchi & Hohwy 2022).

[47] The claim is not that agents and perceivers fully realise the model or that self-aware agents, as a matter of fact, are aware of all the systematic interdependencies that (horizontally) constitute the percepts, intentions, and actions attributable to them. It's an empirical question how far any perceiver or agent approximates the model. Hurley's point, paralleling Dennett and the philosophers of the personal level, is that understanding anything as a perceiver and agent means seeing it as subject to the normative standard set up by this ideal. Below a certain level of approximation to it, it becomes pointless to interpret a system as perceiving or acting in an environment.

entirely disrupted by the effects of SCD. The scotoma is where it is because the rest of the visual system can't make use of the information it would normally receive from these areas.

It's natural to express this by saying that the scotoma 'occludes' whatever happens to be 'behind' it, or that the aura 'takes up' a certain about of visual space. But this elides an important difference in the experience of aura and the perspective a subject has on environmental objects. In ordinary vision, what one sees depends systematically on what one does (and vice versa) in a way that reflects what O'Regan and Noë (2001a, b) call the 'sensorimotor contingencies' of objects in the environment. Roughly, these are patterns of interdependence of perception and action that reflect the presence of an object as a third term mediating their relations: what you see of an object depends not only on what you do, it is a function of what you do given how you are spatiotemporally related to the object. Auras involve no third term of this kind, *at the personal level.* Visual aura percepts 'move with' the head and eye movements that ordinarily bring things into view, and things that would occupy visual space in the area affected by visual aura are either not perceived at all or are subject to perceptual distortions.

Subpersonal explanation of migraine aura in terms of SCD gives us another way to characterise what's going wrong. Looking one explanatory level down, it gives us a way to fill in the missing third term. This involves identifying the mechanisms underlying the perceptual capacities that allow subjects to have a perspective on their environment by taking advantage of these interdependencies. The next step is to locate the phenomenon that is disrupting the functioning of these mechanisms, and so disrupting the corresponding visual capacities of the subject. Doing so allows us to make sense of the pattern of disruption that could otherwise only be descriptively characterised at the personal level. This is the level at which auras shape the experience of their subjects. Let's call this pattern of explanation *interlevel interdependence*:

> *Interlevel interdependence*: Subpersonal mechanism components and their causal organisation constitutively explain personal level phenomena. Mechanism components are individuated in terms of content derived from personal-level interpretation, but mechanism internal causal structure can explain the structure of content attribution at the personal level.

In this case, the interdependencies between the evolution of SCD on the cortex and the aura in the visual field allows for isomorphic mapping between models at distinct levels of explanation: models of visual space, functional models of the retinotopic map in visual cortex, and neuroanatomical models of cortical structure. Since the behaviour of the scotoma is predictable in detail from interactions modelled at lower, subpersonal, levels of organisation, we can

redescribe horizontal constitution conditions of the scotoma in terms of causal and functional relations involved in these models. The availability of explanatory isomorphism indicates that the vertical and horizontal conditions for identifying the phenomenon overlap.

c) Two-level interdependence and perspectival experience

Taken together, *perspectival interpretationism* and *interlevel interdependence* give us the outline of the two-level interdependence view. This can accommodate interlevel isomorphisms in the way just set out. But can we explain states of and capacities for object perception on the same, isomorphic, model as I've suggested we should use to understand the content and phenomenal character of migraine aura episodes? Dennett's early proposals about the subpersonal explanation of states of conscious experience seem to encourage us to do so. These introduce the abstract, functional notion of an 'awareness line' (Dennett 1969, p. 118). Dennett defines an 'awareness lines' for two kind of inner subpersonal mechanisms: those responsible for speech production and those responsible for production and control of directed bodily movement (or 'behaviour'). In each case, the 'awareness line' coincides with the input for these mechanisms, giving us a way to define two corresponding notions of 'awareness':

> A is aware$_1$ that p at time *t* if and only if p is the content of the input state of A's 'speech center' at time *t*
>
> A is aware$_2$ that p at time t if and only if p is the content of an internal event in A at time t that is directing current behaviour (Dennett 1969, p. 118)

As Dennett notes, these definitions "bridge the gap between personal and subpersonal levels of explanation" since the two new terms for 'awareness' "still take persons (or whole systems) as subjects, but have sub-personal criteria" (ibid.). They give us a simple model of isomorphic explanation of ordinary content: at any moment, some combination of the content that an interlevel interpretation gives to mechanisms guiding action and producing speech would count as input to these mechanisms: justifying and explaining the personal level interpretation expressed in terms of a subject's awareness of her environment. So, we can understand the 'awareness line' as an element within a proposed (partial and incomplete) isomorphic model.

This picture is the one Hurley wants us to reject. Ordinary visual experience is of things in the environment. This point is uncontroversial when this is understood as a claim about ordinary causal dependence. But, as we saw in the debate around naïve realism (§1.2), it because more controversial when it becomes about what we can now call *horizontal constitution*. Previous sections have given us reason to think that when constitution is understood in terms of the

contribution these causal dependencies make towards interpretability, we can avoid some of this controversy. As a matter of everyday and scientific explanatory practice, we don't need to make any stronger claims about the ontology of entities at the personal level than we do about social kinds like money or citizenship. While we might have good reasons to do so, these are not immediately relevant to the notion of horizontal explanatory dependence that is involved in the distinction between personal and subpersonal levels.

This kind of horizontal explanatory dependence earns its keep if it allows us to track regularities and distinctions that depend on the relations in which something constituted by these relations participates. In the case of perspectival experience, at least some of the relevant relations are between perceptions intentions and actions. The basic claim is that a subject's visual perspective is constituted by the actual and possible causal interdependencies of these elements. These are part of what someone would need to keep track of to interpret her as a subject and agent. Auras enter the psychological lives of migraineurs negatively, by disrupting ordinary sensorimotor contingencies. Visual space affected by aura is not subject to the same interdependencies between what is seen and done as is unaffected visual space, and so does not support the normal range of action intentions. Visual aura percepts 'move with' the head and eye movements that ordinarily bring things into view, and things that would occupy visual space in the area affected by visual aura are either not perceived at all or are subject to perceptual distortions. So, we can characterise what's going wrong as a disruption of the ordinary dependencies that constitute the subject's visual perspective.

Auras also make a positive contribution to the phenomenal character of experience, accessible through introspection. The language used to describe them attests to the fact that they are introspectable visual objects for their perceivers, occurring in experiences that are best characterised partly in terms of these unusual objects' intrinsic features. If we want to take their isomorphic explanations as a model for subpersonal explanation of perception generally, we have to focus on this kind of positive contribution – it would have to be that everything in the visual field could be accounted for in this way. But this would be to understand the phenomenal features of perspectival episodes on the model of introspection on those episodes – and this is incoherent. Perspectival episodes are constituted by interdependencies between perception and action that are prior to introspection, their phenomenology would not be available as an object of introspection at all if it depended on introspection.

6) Conclusion

The argument just sketched needs to be elaborated and defended. The next chapter will attempt to do this. If this extended argument holds up, then efforts to read the content of episodes of perspectival experience from the contents assigned to subpersonal mechanisms are misguided. The interdependence of action and perception is constitutive of perspectival episodes in a way that can't be captured in isomorphic subpersonal explanation.

If we're convinced by the basic interpretationist insight that questions about what personal-level psychological phenomena are can be usefully approached by asking under what conditions persons are interpretable in terms of those phenomena (e.g., as occupying mental states, being affected by sensations, perceiving objects, etc.) then this chapter makes a prima facie case for the non-isomorphic constitutive explanation of perspectival experience. That is, it makes the case that elements identified in a horizontal constitutive explanation of perceptual experiences structured by a perspective on the environment cannot be put into one-one correspondence with components of neural mechanisms that realise that experience. But saying this need not involve attributing special emergent properties to persons or their perspectival experiences. These might be vertically constituted by the mechanisms that realise them in causal interchange with their environments. Constitutive relevance is strongly suggested by the aura case, and more complex interdependencies emerge as neuroscience reveals more about how the brain realises the mind. Personal-level explanation doesn't rely on a body of knowledge that holds independently of possible further discoveries – it is an interpretive approach that subjects and agents use to understand behaviour. Adopting this strategy or stance guides work in neuroscience and allows for mechanisms to be identified. As they learn more about how these mechanisms function, scientists can revise their original construals of the causal structure behind the behaviour of psychological agents in an environment. Under some circumstances, this allows for an isomorphic explanatory strategy, where the two levels of explanation overlap in models of a given phenomenon. But this occurs against the background of explanations that don't rely on this kind of explanatory isomorphism. And while there are good reasons to think that cognitive scientists rely on personal-level psychological explanation to identify subpersonal mechanisms, philosophers cannot secure a distinct ontology of the personal level by appeal to these explanatory practices alone. Neither are they justified in claiming that these practices are autonomous, and immune from revision in light of a better understanding of the complex causal mechanisms that underly thought, behaviour and experience.

3

Agents, their Actions and the Perspectival Unity of Experience

0) Introduction

Interpretationists think that attending to everyday practices of interpretation and explanation of peoples' reason-guided behaviour reveals the constitutive "order which is there whenever actions are done with intentions" (Anscombe 1957, §42). Extending the interpretationist framework beyond intentional agency means attending to the horizontal constitution of other psychological phenomena, reflected in their position within a network of attributed states, activities and episodes. Since this network is arrayed around what a person says and does, agency plays a constitutive role for the psychological phenomena that are given an interpretationist analysis. For a given phenomenon, the relevant part of this order might be mirrored in the causal structure of neurocomputational mechanisms – but it needn't be. Where it is, isomorphic explanations are appropriate; where it isn't, they are not.

This chapter applies the interpretationist framework to perspectival experience. Perspectival experiences sit at the heart of the networks that interpretationists appeal to. This is because they situate their subject in an environment composed of objects and events, towards which her actions are directed. The capacity for perspectival experience is central to the conception of the person involved in 'personal-level explanation', and acquiring this conception is a key, early step in the process of becoming a subject and an object of psychological interpretation. 'Perspective' and 'perspective taking' are key terms for research into the development of social cognition (Hurley 2008; Moll & Meltzoff 2011). Not coincidentally, perspectival experiences themselves are at the focus of a lot of work in philosophy and cognitive science. 'Visual experience' as it is discussed in classic presentations of naïve realism is a paradigm of perspectival experience, but experiences involving other 'distal objects' such as sounds can usefully be understood as affording their subjects perspectives on an environment made up of what Strawson (1959) calls 'basic particulars'. Developing a neo-Kantian theme from Strawson (1966), Hurley argues that perceiving and acting subjects' perspectives on environmental objects ('basic particulars' in Strawson's sense) determine a core structural feature of consciousness: its diachronic and synchronic unity. The sensorimotor acquaintance view involves a variation on this theme. The appropriateness and availability of non-isomorphic explanation of the perspectival unity of subjects' experiences of environmental objects gives us reason to think that those objects partly constitute the episodes of experience in which they

participate. Since the action-perception interdependencies that constitute subjects' perspectives depend on patterns of causal interaction with the objects of experience, explanation of the content and character of corresponding episodes of experience should advert to properties of those objects.

The subpersonal component of this story shares some of the main elements of sensorimotor enactivism. What explains the constitutive role of the objects of perspectival experience is the integration of sensory and motor processing that causally enacts actual and possible patterns of interaction between subject and object. However, while enactivism has often been presented as an alternative to 'representationalism', the proposed view is compatible with appeals to subpersonal representations, ubiquitous in mainstream cognitive science. Ascribed contents at both levels of explanation are interpretations that situate processes and events (psychological phenomena, subpersonal mechanism components) in their wider environmental, psychological and functional context. Enactivists' rejection of representation, I suggest, is best understood as a rejection of isomorphic representational models, which understand attributed contents as intrinsic properties of vehicles (or of personal-level phenomena). If we take a non-isomorphic, interpretationist perspective, we needn't share the radically revisionary stance that some enactivists take towards representational explanation in cognitive science.

Adopting this approach, though, means ruling out isomorphic explanations of perspectival experience. While it has become increasingly common, in both philosophy and cognitive science to acknowledge some kind of constitutive link between perception and action (see, e.g., essays collected in Engel et al. 2015) – or perception and *action intentions* (Ward et al, 2011; Marchi & Hohwy 2022) – many contemporary proposals are still committed to explanatory isomorphisms. I'll suggest that an important reason for this is the widespread tendency to endorse the standard story of intentional action, in something like the form given to it by Davidson (1980) – but without endorsing the radically non-isomorphic model of propositional attitude psychology that Davidson embedded it in. For Davidson, a *local* correspondence holds between an agent's intention-expressing psychological activity and the 'basic action' causally responsible for realising that intention. Many philosophers of cognitive science assume this local correspondence, which entails that the boundary of the mind and the boundary of an agent's body coincide. This excludes environmental objects, along with everything else external to neurocomputational mechanisms, from subpersonal models advanced on this assumption. In Davidson's case, this is motivated by the rest of his anti-reductionist picture. Since, for Davidson as for the philosophers of the personal level, the explanatory interdependence of psychological events can't be captured in ordinary causal terms, they must have a point of contact with the

environment in order to be integrated into the wider event-causal order. But shorn of these commitments, the standard story can also support a radically isomorphic model. If this is adopted, then the attractive picture sketched above is unavailable. Fortunately, we already have the resources needed to avoid this result. Looking to Anscombe's picture of intentional agency, as developed by recent critics of the standard story (Hornsby 2008, 2013, 2017; Lavin 2013, 2015), we can understand the interdependence of perception and action at the personal level as horizontally constituting an extended, psychological episode in which perceived objects participate. With this in hand, we can call on Hurley's two-level interdependence model to offer a non-isomorphic interlevel explanation, paving the way for the second, positive part of the thesis: the argument for a sensorimotor acquaintance view of perspectival experience.

Here are the steps that take us there. In (§3.1) I discuss and criticise a recent account of perspectival experience shaped by isomorphism assumptions. These encourage us to think that both the content and structure of subjects' perspectives on the environment must be reflected in the content of subpersonal representational states: unified perspectival content is the input to the motor processes that initiate basic action constituting bodily movements. In (§3.2) I introduce the standard story and outline the support it gives to this picture and to isomorphic subpersonal explanation. In (§3.3) I relate this general account back to perspectival experience. In (§3.4) I rehearse Strawson and Hurley's arguments connecting object perception and the subjective unity of perspectival experience, arguing that the standard story fails to support an adequate account of perspectival unity over time. I draw consequences of this for what we can say about the constitutive role of environmental objects in experience.

1) Isomorphic Content and Perspectival Unity

a) Isomorphic explanation and 'Intermediate-level theories of consciousness'

At the end of the last chapter (§2.5c) I sketched an argument that not all subpersonal explanation of perspectival experience can be isomorphic. This relied on two key elements: a) a characterisation of perspective as constituted by action-perception interdependencies; and b) a claim that these interdependencies are the upshot of causal processes that run through the environment. These claims will be crucial premises in my argument for the sensorimotor acquaintance view. So far, they have been stated but not defended. Both are controversial and, moreover, there are arguments that reach conclusions about the content and character of perceptual experience that rely on something like my claim (a) while explicitly rejecting (b). In this section, I'll consider a recent argument along these lines, offered by Francesco Marchi and Jakob Hohwy (2022). Marchi and Hohwy appeal to something like (a) in support of an

'intermediate-level theory of perceptual consciousness' of the kind first proposed by Ray Jakendoff (1987; see also Prinz 2000, 2012). Intermediate-level theories are motivated by and explicitly committed to the kind of general explanatory isomorphisms that I've suggested are incoherent. Since my reasons for this stemmed from a commitment to (a), and Marchi and Hohwy argue from (a) *to* a view tied up with general explanatory isomorphisms, defending my argument requires a response. I'll suggest that reasons for rejecting the intermediate level theory come down to our reasons for accepting (b).

As we saw in (§2.5c), considerations in favour of isomorphic subpersonal explanations of personal-level phenomena can be found in Dennett's (1969) introduction of the personal/subpersonal distinction – the earliest proposed subpersonal explanations of personal level phenomena rely on explanatory isomorphisms. The basic thought is this: if subpersonal explanation of a psychological phenomenon identifies a mechanism or mechanism component causally responsible for events that are interpretable *at the personal level* as an expression of that phenomenon, we can assign the personal-level label to the subpersonal component. Given the moderate interpretationist commitment that interpretation tracks the structure of psychological phenomena, this allows us to appeal to the subpersonal story in a constitutive explanation of the personal-level phenomenon: the mechanism or mechanism components constitutes the phenomenon or capacity.

Although Dennett abandons this account of interlevel explanation in later work (see, e.g., Dennett 1978; Dennett & Kinsbourne 1991), it continues to influence mechanistic and functional explanations of personal-level phenomena. In particular, it influences theories of the content and phenomenal character of conscious states that seek to associate these with the 'subpersonal content' of representational inputs or outputs to putative mechanisms of awareness. Jakendoff's (1987) intermediate-level theory of perceptual consciousness is a theory of this kind. The basics of the intermediate-level theory can be given in three claims (I'll specifically consider the theory as applied to visual experience). First is a claim about how the mechanisms underlying perceptual experience are organised:

> *Hierarchical processing:* Visual perception depends on neurocomputational mechanisms that realise a hierarchy of processing layers. 'Early' layers receive the most immediate sensory input, 'later' layers receive input mediated by earlier layers of processing. These differences are grounded in the functional anatomy of the visual brain, and account for mechanistic decomposition of processing and differences in contents attributable to distinct mechanisms or mechanism components.

Note that the distinction between 'layers' made by *hierarchical processing* is not between explanatory (e.g., personal/subpersonal) or ontological (mechanisms/mechanism components) *levels*. In the terminology introduced by the previous chapter, processing layers are horizontally related by functionally interpretable causal processes. As computational modelling of vision has advanced, the details of Marr's early account – on which Jakendoff's (1987) proposal relies – have been superseded. But as later intermediate level theorists point out (Prinz 2012; Marchi & Hohwy 2022), the mainstream of research on vision continues to support the basics of this picture, according to which representational function is distributed across stages that increasingly integrate features to facilitate object-representation at later and intermediate information-processing layers.[48]

The next two claims are respectively about personal and subpersonal-level content:

> *Subpersonal content:* Contents attributed to 'intermediate level' layers in the functional organisation of visual processing represent objects and events as discrete units, with multiple features as well as more stable properties that survive changes in represented features. These contents correspond to what Marr calls the '2.5-dimensional sketch' of the environment.
>
> *Personal-level content:* Contents attributed to a person's states of visual consciousness represent discrete objects, with transitory features and properties that survive changes in those features.

The natural-looking conclusion that intermediate level theorists reach is that the subpersonal-level contents vertically explain the personal level contents. By understanding how these are sustained and how they interact with other mechanisms in the brain, we get a direct explanatory grasp of (visual) conscious experience. For example, we might understand the intermediate level as constituted by a system of short-term memory buffers that store potentially conscious content from which attentional mechanisms select. In this case, we can construct a theory of awareness that reflects these dynamics. Here's Jakendoff's summary of his version of this proposal:

> The108ction108tionns of form present in each modality of awareness are caused by/supported by/projected from a structure of intermediate level [i.e. an intermediate processing layer] for that modality that is part of the matched set of

[48] I agree with this assessment. Later (§§5.2-5.3) I'll draw on this mainstream work to support the sensorimotor acquaintance view.

> short-term memory representations designated by the selection function and enriched by attentional processing (Jakendoff 1987, p. 49)

As in Dennett (1969), Jakendoff's 'awareness' is an interlevel construct, inheriting its features at the personal level from the dynamics of the mechanisms used to explain it at the subpersonal level. 'Distinctions of form' are distinctions in the features and properties of objects of subjects' awareness, which intermediate level theorists align with contents at the privileged subpersonal processing layer. Jakendoff's open characterisation of the metaphysical relation between levels (i.e., causing, supporting/enabling, or 'projecting') emphasises that this is meant, at least primarily, to be an explanatory or methodological relation: either causal or constitutive relations could underlie the interlevel mappings appealed to on this account – as, in principle, could 'supervenience', understood as covariance of properties and features. The intermediate level theory is a theory of the contents of experience: it offers an answer about what falls within its scope. If contents determine or constrain phenomenal character, it is also a theory of phenomenal character; if the criteria of content attribution track its constitutive conditions, it is a theory of the metaphysical basis of this content and character. But what's essential to the proposal as Jakendoff sets it out is the way it assigns content by appealing to interlevel explanations – and the interlevel explanations it appeals to are isomorphic.

b) Introspection and agency

Above, I described the move from the *personal-level* and *subpersonal content* claims to the intermediate-level theory's explanatory isomorphism between the two as 'natural'; as the sort of thing that easily would suggest itself to anyone who considered the claims that, on my presentation, make up the intermediate-level theory. And I think it is, given the background commitments and theoretical framework introduced by the previous chapter. But the intermediate-level theory requires both claims to be true, and either might be challenged.

Enactivists challenge the subpersonal claim. They argue that the subpersonal representations to which intermediate-level theorists appeal don't really exist. It's a commonplace of cognitive scientific and philosophical discussion that visual experience has (or at least seems to have) particularly *rich* content, where 'richness' corresponds to the amount of semantic content that can be ascribed to a state of experience (Siegel 2011; Hafri & Firestone 2021).[49] Experiences that represent objects, their properties and features, and their respective locations in an

[49] Controversies about visual perception's 'richness' often come down to debates about *just how rich it is* – in this context, the intermediate level theory can be seen as making the case for a *relatively* sparse view (for discussion, see Prinz 2013; Siegel & Byrne 2017)

environment from some particular point of view have, in this sense of the word, richer content than experiences that only represent the presence of a mosaic of sensory qualities. A motivating thought behind one kind of enactivism is that, while personal-level visual experiences are rich, there are reasons to think that no subpersonal representations have correspondingly rich content.[50] From this point of view, since the personal-level claim (or something like it) is true but the subpersonal claim is false, this must mean that "to have the ongoing, occurrent, perception of richness of the visual world, the richness does not actually have to be continually impinging on the retina or activating some internal representation." (O'Regan & Noë 2001b, p. 97). In this case, subpersonal contents cannot isomorphically explain the content of personal-level experience. Since explanatory isomorphisms between subpersonal and personal-level contents aren't available, we must look elsewhere – to patterns of interaction between agent and environment mediated, in part, by these subpersonal mechanisms. [51]

Naïve realists challenge the personal level claim. For naïve realists, perceptual experience doesn't have rich content, either because it doesn't have content at all or because this content isn't essential to perceptual experience (or because perceptual content is itself constituted by relations to environmental objects at the personal level). The upshot of this is that even if the *subpersonal claim* were true, the mechanisms that intermediate-level theorists put at the centre of their accounts would fail to specify and thereby explain the essential, object-involving element of perceptual experience. For this we should look, rather, to the objects themselves – as they are given to the subject in the *sui generis* personal-level acquaintance relation.

Respectively, the naïve realist and enactivist responses just sketched suggest 'top-down' and 'bottom up' strategies for resisting the intermediate-level theory. The moderate interpretationist methodology that I've endorsed above encourages a top-down approach – although one that doesn't rest on naïve realism's challenge to the notion of personal-level content. This says: first assess the grounds for the personal-level claim, then ask whether the reasons for holding it are compatible with the general isomorphism principle that intermediate level theorists rely on. Can a general commitment to isomorphic explanation make sense of all of the features attributed to perceptual experience by the *personal level* claim? If not, we have

[50] For a review of some of the empirical evidence that enactivists have thought supports this claim, see O'Regan & Noë (2001a), Cisek (2007), Menary (2016) and papers collected in Engel et al. (2015)
[51] Radical enactivists (see, e.g., Kirchhoff & Hutto 2016) take this argument further. The contrast they see is not between the rich content of experience and the sparse content of subpersonal representations – the subpersonal level, they argue, contains no representations.

reasons to look to the kind of subpersonal extra-mechanistic relations that enactivists base their approach on.

So, how do intermediate-level theorists support the *personal-level* and *subpersonal content* claims? Marchi and Hohwy (2022) appeal to two sources: the deliveries of introspection and the demands of action. Let's consider these in order. Marchi and Hohwy attribute to Jakendoff the observation that "phenomenal experiences introspectively appear to be integrated [and] subject-centred" (p. 893).[52] Consequently, they think that one reason to endorse the intermediate-level theory is that (given the *subpersonal content* and *hierarchical processing* claims) it makes most sense of these introspectable features of experience:

> [S]ome areas of the brain that are associated with vision display activity that has a sufficient degree of integration, i.e. it is insensitive to minimal shifts in the proximal stimulus (e.g. retinal projection), but is still perspectival, i.e. it is sensitive to changes in subject position. In other words, *the activity of these brain areas appears to have the properties that may be attributed to phenomenal consciousness* from introspection, namely first-person perspective and a certain degree of stability and integration. (ibid.; emphasis added)

On this view, introspection provides criteria for awareness. Subpersonal criteria would have to be provided by subpersonal 'mechanisms of introspection'. The thought that these might be found is behind Dennett's proposal that input to the 'speech centre' of the brain could be appealed to as the subpersonal criterion for 'awareness$_2$' (see, above, §2.5c).[53] But subsequent developments in cognitive science – some of them highlighted by Dennett in later writings (see, e.g., Dennett 1978, chs. 3, 9; Dennett & Kinsbourne 1991) – make this an unpromising option (Schwitzgebel 2008). No uncontroversial candidates are available to provide subpersonal criteria for introspective content.

Moreover, as Marchi and Hohwy note, narrowing down the contents of introspection at the personal level is more difficult than this simple argument for the intermediate-level view allows

[52] They also attribute (and endorse) a third claim, present in the passage from Jakendoff (1987) quted above: perception is 'modality specific'. I don't discuss this claim for two reasons: a) since Marchi and Hohwy focus on vision, the claim is not relevant to their discussion or my critique, b) on the reading intended by Jakendoff and Prinz (2012) 'modality specific' phenomenology restricts phenomenal properties to a narrow range of basic qualities. This claim is controversial as introspective report and implausible as interpretatively attributed content. But while the sensorimotor acquaintance view will ultimately give us reasons to reject this understanding of modality specific content, it would unduly complicate an already complicated discussion to try to show this here.

[53] For some later development of ideas on this theme, see Dennett (2000, 2003)

for. Contents attributed to intermediate-level representations include objects' shape, surface properties and behaviour in an environment, viewed from a particular point of view. Characterising the 2.5 dimensional sketch, Marr assigns it the computational role of "[making] explicit the orientation and rough depth of the visible surfaces, and contours of discontinuities in these quantities in a viewer centered coordinate frame" (1982, p. 37, table 1-1).[54] This role and associated contents attributes both more and less than some philosophers have found, introspectively, in experience. To the extent this contains volumetric information – albeit represented by surface shape 'primitives' (pp. 275-279) – it might be thought to representant more than the sense data appealed to by the early acquaintance theorists. It certainly represents less than the full range of object kind and modal properties that Siegel (2011) motivates via her introspection-based 'method of phenomenal contrast'. And, of course, many naïve realists think that reflection on what is given in introspection gives us reason to deny that perception is essentially contentful.

For these reasons, which "cast doubt on the reliability of introspection in providing the core features that allow [intermediate-level theorists] to pinpoint the intermediate-level as the one relevant for conscious experience" (p. 894), Marchi and Hohwy recognise that whatever support introspection gives to the *personal-level content* claim needs to be supplemented. To do this, they appeal to an 'ecological' consideration,[55] arguing that the features revealed by introspection on perceptual experience – unity and perspectival form – reflect a property of the neurocomputational mechanisms relating perception and action:

> [The] property […] can be shared by some levels of the internal hierarchy and the actions a certain [cognitive] system can perform depending on what kind of system it is (e.g. what kind of body and sensory organs it has). This shared property delimits a specific set of levels in the hierarchy, which, as we argue, happens to be intermediate for humans. (ibid.)

The crucial shared property is supposed to be a range in 'spatiotemporal resolution' appropriate to determining the contents of 'intermediate level' perceptions and a similarly privileged set of

[54] The term is a familiar one in engineering and design. In these contexts, a 2.5 dimensional sketch is one that represents an objects' spatial dimensions on a single plane – contrasted with a 'true' 3 dimensional sketch, which represents all of the values of all of the (represented) spatial features available from any perspective.

[55] 'Ecological' here does not refer to the Gibsonian notion of ecological psychology but rather to the principle, which Marchi and Hohwy take to be derived from ecological science in general, that an animal's capacities are best understood as meeting some survival or reproduction based incentive that that animal has.

'basic' action intentions. By appealing to what they take to be well-established facts about the relation between intentional action, intentions and their subpersonal representation, Marchi and Hohwy aim to "specify how such an intermediate-level may be identified independently from the introspectively detectable phenomenal features of conscious experience" (ibid.).

c) Perspectival experience as an 'action space'

Marchi and Hohwy's proposed agency-based criterion places something like Dennett's (1969) functional 'awareness line' at the input to basic action supporting mechanisms. To understand the notion of a 'basic action' and the role it plays in Marchi and Hohwy's argument, it will be helpful to introduce the supporting notion of an 'action space' (cf. Ward et al. 2012; Cisek 2007. Marchi and Hohwy's ecological argument depends on the claims that the intermediate level determines the contents of the action space and that the action space fixes the content of perceptual experience. Here's a provisional characterisation:

> *Action space (first pass)*: If a set of representations P_1, …, P_n constitutes an action space for a subject, S, then for any perceptually guided action, ϕ^P, performed by S, some P will be part of a reason-giving horizontal explanation for ϕ^P

Action space (first pass) imposes a necessary condition on a set of representations that could be an action space: that its members jointly explain all of the perceptually guided actions a subject makes or could make, and it states a sufficient condition for being an action of S: if some ψ is horizontally explained (in the right way) by P_n then that ψ is a ϕ^P (ie., ψ is a *perceptually guided* action). The relevant notion of 'horizontal explanation' is the combination of causal relevance and functional interpretation that the previous chapter identifies with horizontal explanation at the subpersonal level. It gives us a framework for identifying constitutive elements of a personal-level phenomenon: perceptually guided action.

This provisional characterisation throws up a number of questions. Some of these concern the unity of the representations that are meant to constitute the action space – both at a time and over time. Grouping them into a set for the sake of a definition imposes an artificial unity on the representations, but if the notion of an action space is supposed to give us a better understanding of the relation between perception, action and experience this imposed unity must *in some way* map onto concrete relations at *some* level of explanation. Hurley's (1998) notion of a perspective gives us a sense of where to start at the personal level – she suggests we appeal to relations of dependency between an agent's actions and perceptions as these play out in the environment that agent inhabits (see above, §2.5b). But so far it is unclear how these are supposed to relate to underlying subpersonal mechanisms. And, as we'll see, the notion of an

action space as Marchi and Hohwy use it is meant to cut out these concrete subject-object relations. For them, unity can be understood in terms of the contents of the representations that make up the action space. I think there are compelling reasons to reject this solution, which I'll lay out when I return to these issues below (§3.4).

Another worry, which Dennett (1969) voices in his discussion of the related notion of awareness$_2$, is that the requirements of 'current behaviour' fail to plausibly constrain the contents of a set of representations invoked to explain it. To put the worry in terms of the scope question alluded to above: if the notion of action guidance is supposed to constrain possible contents of perceptual experience, we have no reason to narrow down these contents in the way supporters of the intermediate theory want us to. Asked why Homer threw the ball at the snowman (see above, §2.2b), we *might* say that Homer on one hand perceived an object (the snowman) at such-and-such a distance with some set of 'intermediate-level' properties (whiteness, roundness, etc.), and on the other hand stood in the right kind of non-perceptual psychological relations (to the snowman, to Ned, etc.) that he was guided by this perception to the resulting action. But we do not yet have a reason to prefer this over other introspectively at least as plausible alternatives that attribute more (or less) semantic content to Homer's perception.

This is where appealing to the spatiotemporal resolution of basic actions is supposed to help us. Consider three descriptions of Homer's action in throwing the ball:

a) Homer activates muscles, XYZ, causing his ball-throwing hand to describe a path between spatial points, α, β, before releasing it at point β at a time, t
b) Homer throws the ball at a time, t'
c) Homer accidentally breaks Ned's window at t'_{+1} by throwing the ball at t'

Of these, (a) has the most, and (c) the least, precise spatiotemporal resolution. (a) specifies a series of events inside Homer's body that occur over a short interval terminated by the instant, t, at which Homer's fingers release the ball he is throwing (or, if you prefer, over the *very* short interval in which his fingers were releasing the ball). (b) involves all of these events (and so the interval of time that they occupy) but extends the frame of description beyond Homer's body to include the ball and – at least plausibly – beyond the interval terminating at t. In order for the act to be a throwing, we might say, the ball has to be – or have been – *thrown*. (c) encompasses the interval t' as well as the events that occurred after Homer released the ball, hitting and breaking the window. Let's assume that representations corresponding to (a) and (b) are realised at some level of the neurocomputational mechanisms underlying Homer's agency –

following Marchi and Hohwy, we can call representations of this kind 'control states' for the prospective actions they represent. In this sense, no content associated with (c) can be a control state. (c) dramatizes the fact that intentions can fail to produce intended distal outcomes, but assuming that there are neurocomputational mechanisms whose function is best understood as representing *target* outcomes, we can assign contents of the kind illustrated by (c*) to them:

c*) Homer intentionally knocks over Ned's snowman at t'_{+1}

Of course, if (a) (b) and (c*) are to be plausible contents for *Homer's* representational states they won't have the third-personal mode of presentation that I've given them here. But since specifying this aspect of their content, even at the personal level, raises a further set of difficult questions, I'll set this part of the problem aside.[56]

Recall that to interpret Homer as knocking over Ned's snowman, we have to take into account more of the context of Homer's action. I'll now make some of that context explicit. Homer intends to knock over Ned's snowman to avenge himself on Ned for imagined slights. Here we might picture a thought-bubble over Homer's head, depicting a series of events: first the throwing of the ball; then the snowman's tumbling over; then, for instance, Homer gloating over a suitably humbled-looking Ned. The contents of the bubble correspond roughly to what Marchi and Hohwy describe as a 'policy'. A policy is a complex representation that links a series of possible actions to the goal(s) of an agent. Policies do not determine the particular actions that realise them – typically, a number of different possible actions would count as realising the same policy. They have correspondingly low spatial and temporal resolution (by analogy: think of all the different ways in which the contents of Homer's thought bubble might be depicted).

Nearly everyone who makes a distinction between personal and subpersonal levels agrees that the highly precise contents associated with (a) are subpersonal.[57] Homer's perceptions do not (typically) give him reasons to perform actions described this way, and we need to introduce artificial terms to specify what ordinary, personal-level action description leaves unspecified. Representations like this are not candidates for what fixes the content of Homer's perceptual experience. Marchi and Hohwy's diagnosis for this is that the intentions with which Homer acts

[56] I take it that there is some way questions about the reflexive content of intentions could be satisfyingly answered, and so Marchi and Hohwy can help themselves to the this account – and that the same holds for sensorimotor acquaintance theorists. If no good answers are available, then they all inherit this problem together.

[57] For a dissenting discussion of a similar case, see Bermúdez (2000, pp. 79-80). The conclusion Bermúdez reaches, that 'subpersonal' actions like this are in fact attributable to the agent is, I think, in line with the conclusion I end up reaching – the personal/subpersonal level, as I understand it, is not incompatible with interlevel identities.

are indeterminate with respect to the precise way in which an action is performed. But, they continue, Homer can't act on his policies without performing some particular (basic) action. Here the situation is reversed, in order to realise the policy Homer must act in some specific way. From this, they conclude that:

> [i]t is reasonable [..] to associate basic actions with control states and complex actions with policies. Executing a complex action is then inferring a policy, and executing basic actions is then realising each control state that makes up that policy. (p. 901)

So, (b) type contents are tied up with the most basic actions that Homer performs. If perception is supposed to provide Homer with an action space, then it must at least provide him with the content he needs to guide his basic actions. But granted these, he seemingly has *everything* he needs to act effectively on his policy:

> Control states can then be thought of as specific causes of movement that make the complex action tractable given the system's capacities and body, by breaking down the complex action's spatiotemporally distal goal states into smaller steps with more egocentrically proximal goal states (ibid.)

Since actions are now tractable, we can understand what Homer does as the serial execution of a policy and the causal upshot of the mechanisms that bring about, moment-by-moment, the more precisely described bodily movements by which he is causally effective in the world.

But if this is all we need for a subpersonal explanation of Homer's acting the way he did – and so all we need to characterise the action space that, according to Marchi and Hohwy constrains the content of Homer's perceptual experience, then it looks like we can provide a generally isomorphic subpersonal explanation. This understands the relationship between what is perceived and what is done in terms of the relationship between control states, basic actions and their perceptual consequences. Together these map out the resources on which a vertical, constitutive explanation could draw. And in this case, objects are constitutively excluded from perspectival experience. This means that, if we want to take advantage of the theoretical benefits of acquaintance (§1.1), we cannot understand the relationship between personal and subpersonal level in terms of constitution.

2) The Standard Story of Action and Isomorphic Explanation

a) Introducing the standard story

In action theory, the notion of a 'basic action' is frequently associated with the standard story of intentional action and Davidson's (1980) work elaborating and defending this story (see, e.g., Hornsby 1980; Velleman 1992). While Marchi and Hohwy distinguish 'basic actions' in their sense from those that are at issue in the debate about 'metaphysical' basicness (Marchi & Hohwy 2022, p. 899), it is far from obvious that they can hold considerations raised by the metaphysical debate apart from what they describe as the "intuitive characterisation of a complex action" as one which is "composed of a series of simpler building blocks, namely a series of basic actions" (ibid.). I think that this picture is quite often the one that action theorists are debating. That it seems intuitive to many philosophers of mind and cognitive science whose focus is on other questions, might be seen to be a consequence of the widespread adoption of a similar set of assumptions to those at the heart of the standard story (for discussion, see Kalis 2019).

According to its advocates, these assumptions are very deeply rooted. Davidson motivates the standard story by stressing the straightforwardness and plausibility of "the ancient – and common-sense – position that rationalisation is a species of causal explanation" (Davidson 1980, p. 3).[58] 'Rationalisation' here refers to reason-giving psychological explanation: rationalisations explain the behaviour of agents in terms of their motivating reasons, and this behaviour is made up of intentional actions to the extent that it is made intelligible by reasons impugned to the agent. That is, an action is intentional just in case it is done for some reason. This gives us a general way of understanding intentional actions and their explanation:

> Whenever someone does something for a reason […] he can be characterized as (a) having some sort of pro attitude towards actions of a certain kind, and (b) believing (or knowing, perceiving, noticing, remembering) that his action is of this kind (pp. 3-4)

What falls under (a) is paradigmatically a 'wanting' or a 'desire' but is meant to include any evaluation the agent might make of the prospective action as worth doing. In Anscombe's terms, which Davidson adopts, they are evaluations of an action as having a "desirability characteristic"

[58] Whether the assumptions behind the standard story can – or can uniquely – be traced back *this* far is controversial (Lavin 2013; Hornsby 2008). Hornsby suggests that the story is "facet of a much more widespread tendency to think of what there is atomistically" (2013, p. 2) and elsewhere (2008) associates it with the same kind of Cartesian assumptions that Dennett and Hurley associate with explanatory isomorphisms.

(Anscombe 1957; Davidson 1980, ch. 1). When (a) and (b) come together the agent has a "primary reason" to perform the relevant action.[59]

Considered in themselves, movements of the body typically don't have desirability characteristics. Davidson, like Marchi and Hohwy, is impressed by the fact that most of what agents intentionally bring about are the effects of movements of their bodies. In Davidson's now-classic example: someone's understanding that flipping a switch will cause a light to go on, and so illuminate the room, is (in combination with her other beliefs and goals) what makes moving so as to flip the light switch interpretable as an action of the intended kind. At the same time, ascribing this belief-desire pair (or something similar) to her allows an observer to make sense of her flipping the switch. But when someone does this, given how things are with the world, she doesn't perform two distinct actions: first flipping the switch *and then* illuminating the room. Rather, flipping the switch in these circumstances *just is* illuminating the room, though it's only under the latter description that the action has the desirability characteristic needed for *this* action to be rational (and so subject to rationalising explanation). This claim has since become known as the Anscombe-Davidson principle, and is supposed to account for how actions, which are bodily movements, are at the same time realisations of agents' intentions in acting.[60] In the terms introduced by the previous chapter, it supports horizontal constitutive explanations of such facts.

But from this point, Davidsons's account diverges from Anscombe. Recall that, for Anscombe, the relations between an agent and her action are not specified in causal terms. I've proposed the label 'horizontal constitution' for the way that Anscombe, and the other philosophers of the personal level, saw psychological phenomena as interdependently related to each other in a network whose structure is revealed (and relied upon) by the practice of interpretation. For Davidson, intentional action constitutes a point of causal contact between networks of this kind and an environment governed by ordinary event-causation. Like Dennett's, his account is meant

[59] In Davidson's formulation, desirability characteristics are attached to actions via their featuring in propositions that ascribe these to the action – e.g. 'I want *that* I ϕ', 'S desire *that* S be ψ-ing'. As Hornsby (2013, 2017) has pointed out, this leads to problems if the content attribution is interpreted realistically, as the intrinsic content of a psychological state. While I don't pursue the issue, the sensorimotor acquaintance view is compatible with thinking that objects of intentions aren't – or needn't be – propositional contents.

[60] The Anscombe-Davidson principle is accepted by many contemporary action theorists and is implicit in the model of the personal subpersonal account that I endorse. But it also has significant critics (e.g., Alvarez & Hyman 1998). I think part of what motivates these critiques are the consequences it has embedded in Davidson's version of the standard story or – maybe even more so – when it is advanced as part of a general isomorphic explanatory strategy. So, we might understand the chapter (§2) as trying to detach the principle from these problematic associations.

fit intentions and intentionality into a broader scientific picture of the causal structure of the world – but the point at which the world's causal order impinges on psychology is situated on the same, personal-level, of explanation as the psychological phenomena that make up the interpretational network (for discussion, see Dennett 1991b, Child 1994, ch. 1; Brook 2013).

Davidson's proposal is that the explanatory connection between an agent's primary reasons and her intentional actions faithfully tracks a causal connection between psychological events (i.e., the onset of some element of the agent's primary reason), movements in the agent's body that these cause, and the further worldly effects of these movements. Here's how it works. Given that an agent wants to illuminate the room and understands that flipping the light switch will bring this about, she has a primary reason to move her body in order to flip the switch, in order to illuminate the room. The primary reason, whose content reflects this means-end structure, is understood as something internal to the agent who is the causal source of what happens in the world (the agent's finger presses down on the switch, which snaps the switch gate closed, which completes the circuit, etc… *which illuminates the room*). This is what brings these intended consequences under the agent's self-directed causal power; she is the agent of the distal events that reflect the content of the intentions effective in bringing them about. Causally efficient primary reasons "constitute certain events free and intentional actions" (p. 19). Free, because the119ctionn originates in the agent's own psychological states; intentional, because they are, under some description, motivated by reasons the agent endorses.

Davidson is concerned to map the means-end structure of practical reason onto a discrete causal chain originating in the agent because of his dual commitments – to realism about the psychological kinds that make up primary reasons and to a scientifically naturalistic picture of the causal relations between agents and their environments. Agents and their intentional actions must be fitted into the natural order as described by science; and for Davidson, causal connection between events is at the centre of this order. Event causation is "the cement of the universe" and our understanding of these causal relations "is what holds together our picture of the universe … [which] would otherwise disintegrate into a diptych of the mental and the physical" (Davidson 1980, p. xi). Davidson puts forward the causal theory in order to respect the intuition that there must be a point of connection between the causal order of the world and the rational normativity at the centre of our folk-psychological explanatory practices, which he takes to be embodied in both common sense and science.

b) Davidson on primary reasons and basic action

If a naturalistically acceptable account of how agents relate to their bodies and their environment requires that rationalisation track causation in the way that Davidson proposes, then objects cannot feature in (naturalistic) constitutive explanations of the phenomenal character of perspectival experience. The reason for this can be expressed in a slogan from a later paper also collected in Davidson (1980). If the standard story is true, it follows that "[w]e never do more than move our bodies: the rest is up to nature" (1980, p. 59). While agents can be responsible for the distal effects of their actions, these are 'up to nature' because after the agent moves her body, she has no further contribution to make to the actualisation of her goals. The world must cooperate; but even if it does, she has at least done *something*. For Davidson, this means that all that the agent does comes down the bodily movements involved in causal chains initiated by the agent's primary reasons. Call this the *strict bodily constraint* on action.

Davidson sees the strict bodily constraint as following from the Anscombe-Davidson principle and the causal account just outlined. I'll first informally set out this argument before extracting what I think are its key elements below. One advantage to understanding the intention-action relation as causal is that it suggests an answer to a puzzle about how primary reasons specify the events that satisfy them. The answer follows a similar pattern of reasoning to the one followed by Marchi and Hohwy (2022). Causation is a relation between two or more events (e.g. a bodily movement and its psychological cause), but if the psychological particulars that make up a basic reason weren't linked causally to the events in the world that they rationalise, how could one rationalise the other? A natural thought is that the content of the belief-desire pair picks out the *particular* bodily action that is to be performed. But Davidson rejects these on the grounds set out above (§3.1c):

> [T]he event whose occurrence makes 'I turned on the light' true cannot be called the object, however intentional, of 'I wanted to turn on the light'. If I turned on the light, then I must have done it at a precise moment, in a particular way—every detail is fixed. But it makes no sense to demand that my want be directed to an action performed at any one moment or done in some unique manner. Any one of an indefinitely large number of actions would satisfy the want and can be considered equally eligible as its object. (Davidson 1980, p. 6)

Because it specifies a type of action that might be performed in different ways, the contents attributed to agents in rationalising explanations are indeterminate between many candidate bodily actions, which count equally as instances of the target action type. While subpersonal 'control states' that mediate between high-level representations and motor output fix the

details of the bodily movement much more precisely, this is not content that we attribute to the agent. But this disconnect is unproblematic if the vehicles of personal-level intentions are causes of this lower-level processing, since these are guaranteed to link up to the particular environmental events that bring about an agent's intentions. With this causal connection in place, rationalising explanations can exploit it, tracing the means-end structure of the performed action back along this causal chain.

Davidson (1980, ch. 2) fills in some more detail on how this model leads to the strict bodily constraint. Tracing the means-end structure of a rationalising explanation along the event causal chain brings us to an action that is a means to an end (e.g., the light turning on) without itself being the end of some other intentional action. This is an agent's direct causal intervention on the environment, rather than an event further along a causal chain that she is the author of:

> Not every event we attribute to an agent [i.e., ordinary actions like turning on a light] can be explained as caused by another event of which he is the agent: some acts must be primitive in the sense that they cannot be analysed in terms of their causal relations to the same agent. (Davidson 1980, p. 49)

These primitive, or basic,[61] actions are always movements of the agent's body that are (at the personal level of description) proximally caused by her occurrent primary reasons. Importantly, basic actions are not causes of ordinary actions (though they do cause events that occur in descriptions of them). Flipping the switch and illuminating the room are one and the same action *de re,* (i.e., a particular movement of the body) though they fall under different descriptions. Since flipping the light switch is a way of illuminating the room, the intention to illuminate the room is realised by this bodily movement. But as flipping the switch (moving the body so as to do this) and illuminating the room are in fact the same action, the action-constituting event can be precisely located in just these bodily movements as caused by the relevant intention.

c) The strict bodily constraint and neurocomputational mechanisms

We can reconstruct Davidson's argument for the strict bodily constraint to look like this:

[1] An agent's actions are caused and rationalised by her primary reasons.

[2] Actions are rationalised in light of their distal effects.

[3] Primary reasons directly cause movements of an agent's body.

[4] Movements of an agent's body cause distal effects.

[61] Danto (1965) introduces the term, 'basic action' which has become standard in the literature (Amaya 2017).

[1 – 4] state the premises of the argument for the strict bodily constraint. [1] is a statement of the standard story's guiding idea; granted this, [2 – 4] look like uncontroversial claims about the psychology of action and the causal relations between an agent's inner states, events in her body and events in the environment. From these, we can see how the standard story is supposed to solve the puzzle about bodily actions and the indeterminate content of intentions:

[5] From [3] and [4], primary reasons indirectly cause distal effects via bodily movements

[6] From [1], [2] and [5], some movements of an agent's body are actions

[5] accounts for how bodily movements are able to meet the causal criterion for actions, since they are caused by primary reasons and cause distal goals to be met; [6] establishes that these movements of the agent's body also meet the rational criterion for actions. Since other means for primary reasons to cause distal goals are excluded, this gives us:

[7] From [5] and [6], there are no actions that are not also movements of an agent's body

Passing over worries about intentional omissions, looks unproblematic. Now comes Davidson's (1980, chs. 8-9) proposal about event individuation: events that share all their causes and effects are identical. This isn't an entirely uncontroversial metaphysical claim, but it looks hard to resist when differences in either cause or effect is understood as a sufficiency condition for non-identity between events – and this is all that is needed for the argument to go forward.

[8] If two events, A and B, have distinct causal relations, then A and B are distinct events

[9] From [3], [5] and [8], bodily events and their distal consequences are distinct

But now introducing the Anscombe-Davidson principle seems to lead directly to the strict bodily constraint:

[10] For two actions A, B, if A is rationalised as a means to B, then A and B are the same action

[11] From [8], [9] and [10] all actions are only changes in an agent's body

And if all actions are only changes in an agent's body, then whatever goes on outside the agent is not, strictly speaking, part of her action.

I think this argument can be resisted in at least three ways: a) by rejecting the Anscombe-Davidson principle (i.e., rejecting [10]); b) by rejecting the identification of actions with events

(blocking the move from [5] to [6]); c) by rejecting the causal link between reasons and actions (rejecting [1]). Below (§3.3a), I'm going to present an argument that takes the third option. This approach is influenced by the recent tradition of criticism of the Davidsonian standard story. These proposals, however, often buttress their rejection of the causal model with a critique of the Anscombe-Davidson principle or of the event-causal framework in which Davidson situates human agency. But, as I understand it, the Anscombe-Davidson principle is a special case of the more general principles behind the model of interlevel relations developed in the previous chapter. While some of the contemporary critics of the standard story rely on additional considerations about which the sensorimotor acquaintance view is neutral, the Anscombe-Davidson principle is a crucial part of the conceptual framework on which it relies. So, without knocking away this essential piece of supporting structure or appealing to an expanded ontology of personal-level processes and activities, the plan is to show that the isomorphism assumptions that Marchi and Hohwy share with other advocates of the standard story fail to make sense of the features of perceptual experience they seek to explain.

First, let's review their argument. Plugging the conclusion in favour of the strict bodily constraint back into Marchi and Hohwy's case for the intermediate level theory, we can now see why we might expect an action space for basic action to be exclusively constituted by the contents of neurocomputational mechanisms. If intentions and actions are causally linked in this way, then basic actions are the constitutive outer limit of the mind. The point at which the interpretable, intelligible states, events and episodes meet the causal order of the environment. As Doug Lavin (2013, p. 274) notes, agency so conceived, is "a sort of metaphysical containment wall" between psychology and the causal order of the environment. It excludes from constitutive relevance to each other the phenomena that fall on either side of it. So, when we ask what constitutes an agent's intentional agency, we have only internal resources to call on – the immediate causal history of her basic actions. But while Davidson resists the idea that contents at the personal level can be associated with and explained by neurocomputational mechanisms (Davidson 1980, ch. 11, 13), theorists unconvinced by this anti-reductionist stance should find it natural to advert to mechanisms and their components, identified along the lines set out in the previous chapter (§2.3-2.4).

But why in particular should we expect *intermediate level* contents and basic actions to be linked? The reason for this comes from the *subpersonal content* claim: intermediate level contents are supposed to uniquely display the integration and perspectival structure that (basic) actions rely on. In order to act on her environment, a subject must represent the objects and events that it contains together, sharing a common spatial frame of reference with each other

and with her own body. Representing more than this is helpful for framing policies (Homer represents the snowman as belonging to Ned, Davidson's agent represents the light switch as a potential cause of illumination), but these 'high level' properties are not immediately relevant to the basic acts by which the policies get realised. The reason for this, Marchi and Hohwy suggest, is that these properties track changes and relations that the objects participate at a lower degree of spatiotemporal resolution. But, to ward off potential counterexamples (e.g., some causal disposition might covary with an objects' basic perceptible features and properties), we could make a weaker claim to the same effect: there is no essential link between the temporal scale on which these high-level properties evolve and the temporal scale on which the agent – basically – acts. If all an agent intentionally does can be understood as basic actions and their concatenation, then the spatiotemporal resolution of these bodily movements would have a privileged role in determining the representations that constitute her action space.

3) Perspectival Unity and Basic Actions

a) The basic action space model

In order to follow through on the plan sketched above, let's consider the *personal-level content* claim in some more detail. This states that perceptual experience is perspectival and unified. In light of Marchi and Hohwy's ecological argument, we can now say that it represents a set of objects together in an action space that play a privileged role in guiding a subject's basic actions. Part of what's behind this argument is the thought that perceptual experience is rendered unitary by the requirements of agency. For Marchi and Hohwy, this unity is part of the content of an experiential state – it is an upshot of the fact that experiences have complex content that might also be represented as a conjunction of simpler contents. Below (in §3.4), I'll give reasons to reject this content account of unity and prefer a structural account, on which distinct contentual states are seen as contributing to perspectival unity. But for now, the concern is not with how unity is realised – by one overall state or by multiple distinct states – but with what it allows creatures that enjoy unified experiences to do. Experience could not provide grounds for action if it did not give the agent a unified perspective on the environment around her. Conversely, if the agent were disunified – if her motor and cognitive capacities were not functionally integrated in a way that supported effective action – then a unified perceptual field would be of no use. This provides ecological reasons to expect experience to be unified – but these go by deeper ecological reasons to expect that the agent herself be a unified source of causal influence on the environment.

This means that unity and perspectival structure are contingent aspects of the content of perceptual experience. Marchi and Hohwy's argument aims to show that attributions of unified perspectival structure, "while not being an essential feature of phenomenal consciousness itself, may be contingently true of some cognitive systems, given the constraints on their active interactions with the environment" (p. 894). For agents like us, the thought is, these constraints are provided by our repertoire of basic actions. This overlays a model of the structure of intentional action based on the standard story onto the ecological considerations that motivated us to understand perspectival experience as an action space. The unity of the contents that make up the action space comes down to the fact that any of them might serve as an input to the control states that guide basic action, and what links these to intermediate level content in particular is that these represent the environment at the appropriate degree of spatiotemporal resolution:

> *Basic action space*: If set of representations $P_1 \ldots P_n$ is an action space for a subject S then for every perceptually guided basic action, ϕ^P, performed by S, ϕ^P some P is an input to a control state, C^ϕ, that causes and rationalises ϕ^P

In order for this proposal to work, the contents of the control states that cause basic actions must be attributable to the agent who, in performing them acts. The way the standard story integrates the event causal and psychological orders requires that explanatory isomorphism holds. But some complaints made by many of the standard story's contemporary critics give us reason to deny that this isomorphic interlevel relation gives us an adequate interpretive grasp of human agency. The kind of contents that this story attributes to causally efficient control states fail to rationalise the actions they bring about. Notably, these arguments are couched in a similar set of worries about *agential* alienation to those raised by acquaintance theorists about the alienation of the perceiving subject from her environment.

b) Alienated again?

A major theme in recent critical discussion of the standard story (Hornsby 2008, 2013, 2017; Lavin 2013, 2015; Small 2016) is that this way of welding the agent and the causal order together somehow alienates an agent from her own actions. Lavin, for example, compares the position of the agent so-conceived as "not different in principle from Marx's understanding of the relation of the non-worker (*Nichtarbeiter*) to the material processes that realize his own ideas […] Each is alienated from the progress, or getting done, of his deeds."(2013, p. 274). Hornsby, meanwhile, suggests that the causal-compositional account

> encourages a sense of alienation by speaking as if you stood to an event that is your action in a relation expressible using the word 'do'. This makes it seem as if you could participate as agent only by being related to something that might be present in a scene in which you yourself were not involved. (2008, p. 174)

Both Lavin and Hornsby evoke the notion of the 'process' or 'activity' involved in a agents getting something done I take it that this involves the notion of *ongoing* action (see, e.g. Hornsby 2017) – that the alienation that Hornsby and Lavin are concerned with is alienation from the temporally extended activity of getting things done.[62] This temporal aspect is crucial to the case I want to make below (§3.4). But even for those who accept Davidson's event-causal framing, there is something wrong with the standard story as it's been presented so far. Along these lines, David Velleman (1992) claims that the account as stated is only adequate for describing actions "from which the distinctively human feature is missing", telling us "not what happens when someone acts, but what happens when someone acts halfheartedly, or unwittingly, or in some equally defective way" (p. 462). These authors all agree that there is something wrong with the standard story and diagnose the problem as a kind of absence or defect. Even if it succeeds in integrating 'actions' of some kind into the naturalistically described causal order, the view somehow leaves *agents* out of the picture.

This line of objection contests the claim that bodily and worldly events that constitute complete actions can be specified independently of the agent and her intention (e.g. as the relation between a cause and an effect). It's this claim that gives rise to the strict bodily constraint on action, because if mind and world are to be kept apart, there must be a discrete point at which they interact: this is the bodily action, caused by the intention-constituting control state. If we only think about cases like switching on a light, this might not look like a problem. But now consider temporally extended actions, in which worldly, bodily, and neural events and structures mutually interact. Davidson (1980) gives a straightforward example:

> When I tie my shoelaces, there is on one hand the movement of my fingers, and on the other hand the movement of the laces. But is it possible to separate these events by calling the first, alone, my action? (p. 51)

[62] Davidson (1980) is a landmark text in a tradition of discussion of the ontology of processes and events and action that raises doubts about the notion of an 'ongoing event' (for discussion, see Bacharach 2021). Nearby considerations are at work in Hornsby's criticism of the standard story, but I will not attempt to engage with these here. For now, it might be worth clarifying that the notions of an 'ongoing action' and 'temporally extended activity' introduced in the main text are not meant to be 'perfective' in the way some authors have found actions and events to be.

Davidson answers 'yes' to his own question,[63] even while noting that the bodily movements involved in shoe-tying might not be available to the agent without feedback from the shoes and laces themselves, and are not (as, indeed, with the light switch example) directly the object of the relevant intention. This is because they are still intentional under some description, and (once the subpersonal implementational details are filled in) are caused by a content-bearing state matching this description. What this means is that for each bodily movement involved in tying the shoelace, we have to posit a causal link to an intention in the part of the agent. But the content of this intention, at any moment, cannot be to move her fingers just in the way they are moved. This is clear in that one cannot make the required movements on command.[64] Davidson resists the natural thought that the basic action(s) of (involved in) someone's tying her shoes are an event that subsumes the movement of the laces because this would displace the agent and her intention as the unique cause of the action-constituting event. This must be what the agent does "with his body alone […] separate from whatever else takes place" (ibid.), or else we have to see what happens as an event somehow co-authored by the agent and her shoelaces. But, in avoiding this seeming absurdity, [65] Davidson sets up the conditions for agential alienation:

> [A]n agent always knows how he moves his body when, in acting intentionally, he moves his body, in the sense that there is some description of the movement under which he knows he makes it. Such descriptions are, to be sure, apt to be trivial and unrevealing; this is what ensures their existence. So, if I tie my shoelaces, here is a description of my movements: I move my body in just the way required to tie my shoelaces. (ibid.)

It's possible to make out two interrelated kinds of alienation here, epistemic and metaphysical:

Epistemic alienation: This picture seems to involve a disconnect between the know-how of a skilled agent and her manifestation of it in her activity. On Davidson's account, that an agent

---

[63] The question here being *is the action* just *a bodily movement with the right kind of mental cause?* There is a related question about whether descriptions under which basic actions are intentional can refer to their extra-bodily effects. On my reading, Davidson thinks they can, but that the event that realises (and is caused by) the intention will still be the bodily movement only (Davidson 1980, pp. 51-52). The contents of personal-level intentions needn't, and generally don't, specify the determinate bodily actions that satisfy them.

[64] "[T]his is a trick I might learn" (Davidson 1980, p. 51); but if I have not learned it, it is not part of the repertoire of actions I can do *just like that* under any circumstances. This might mean that shoelace tying is basic under some circumstances and not others. Presumably, if I try to learn to perform the trick of moving my hands just so, unaided by the laces, I will do this by doing other things at least while I am still practicing.

[65] Note that this appears absurd *if we assume that being the cause of getting the laces tied is tied up with being the intentional agent of the tying* – since, of course, shoelaces aren't intentional agents they can't be agency-constituting causes.

doesn't know how to move so as to tie her laces outside of the context of ordinary performance is given as a reason to think that she lacks this knowledge (at the personal level) when she has the laces at the tips of her fingers. Rather, only a bare common factor (a high-level motor intention, an internal representational state) is supposed to account for what's essential to the phenomenon under analysis.[66] Critics of the standard story urge that characterising skilled activity this way distorts it, since bodily skill is always manifested in context. Classic examples in the causal-compositional account typically involve unskilled action (pressing a button, raising one's arm) but it's worth remembering how unrepresentitive actions of this kind are. If ordinary action is the deployment of some skill or capacity, then we might worry that it is mischaracterised by this approach. Ordinary agents, we might think, exhibit a more intimate understanding of what they're doing than the Davidsonian view accounts for.

Metaphysical alienation: This epistemic worry leads onto another, more metaphysical worry. If we are tempted to think of agency of as the exercise of agential skills or capacities, we might wonder about how this is related to the intention that is supposed to be doing the causal work. We attribute skill to the agent as a whole rather than her causally efficacious mental states (we speak of a skilful agent rather than a skilful intention; cf. Steward 2000). But it's not clear how such capacities are manifested in punctuate mental events that are at the centre of the Davidsonian picture. That is, we might wonder how the agent, rather than some subset of mental events within her, gets to be efficacious. So far, we have assumed that the integration of motor, perceptual, and cognitive resources in determining a unified action is enough to ground synchronic experiential unity. But, at any such instant of agential intervention, we might ask, is the agent herself really present?

c) Constitutive rationality and perspectival unity

Let's take these concerns in turn. The epistemic worry raises a problem for defences of intermediate level theory that appeal to the basic action space model, that can be expressed as a dilemma. Either:

[66] This reasoning fits into a pattern of other 'common-factor' arguments about perception and knowledge: a bad (decontextualised) case is used to demonstrate that something seemingly present in the good (contextualised) case really isn't. Rather, only a bare common factor (a high-level motor intention, an internal representational state) is supposed to account for what's essential to the phenomenon being investigated. Recall that common factor arguments about the phenomenological, epistemic and semantic properties of perceptual experience are what many arguments for naïve realism are supposed to enable us to resist. Drayson (2018) notices something similar going on in the debate around embodied cognition.

a) The control states that manifest the agent's knowledge how to perform a skilled action are personal-level intentions that determine the contents of her basic action space

or

b) these control states are subpersonal, they do not determine the contents of her basic action space.

(b) is the option suggested by Davidson, but for him the standard story is not part of a wider strategy for linking the contents of perception to a specific set of representational states. In the context of this argument, the disconnect between what an agent could, basically, intend to do and the range of skilled actions she can perform threatens to undercut the link that Marchi and Hohwy want to maintain between the standard story and the ecological principles that they use to motivate the action space model. But the other horn of the dilemma, (a), threatens to break the link between basic intentions and intermediate level contents. It seems plausible that representations both with more and less spatiotemporally precise contents would be plausible control states for skilled actions like these. On one hand, fine details of action execution depend on transitory features of a perceived scene; on the other, tracing causal dispositions and modal features seems essential to many kinds of skilled performance.[67]

While these considerations put pressure on Marchi and Hohwy's (2022) argument, I don't think they are decisive. We can nonetheless leave the debate here. I introduced their argument for the intermediate level theory to illustrate how someone committed to the explanatory isomorphism principle might accommodate the idea that perception and action are interdependent. We've now revealed some costs to this approach: both those involved in specifically agential alienation and those involved in the denial of acquaintance explored above in (§1). To return to the present line of argument, imagine someone who takes the first horn (a) of the dilemma and the consequences that I've suggested follow from it. This person would think that a link between perception and action intentions determines the content of perspectival experience, and that this content can be associated with the contents of a set of neurocomputational mechanisms underlying action guidance.

How would someone like this respond to the metaphysical worry? One way forward is suggested by Velleman (1992). Velleman's rejection of Davidsonian primary reasons is partly motivated by this kind of concern, which he names the "problem of agency". The problem of agency is the

[67] As Hurley (1992, 2008) and more recent theorists (e.g. Hesp et al. 2019; Ramstead et al. 2020) have argued, the functional unity of living systems does not depend on causal bottlenecks of the kind that Velleman and Davidson present as necessary for human action.

problem of locating the agent in the economy of mental events (desires, beliefs, perceptions, intentions) appealed to on the standard story. Recognising this problem leads some philosophers to propose a *sui generis* kind of agent causation (see, e.g., Bishop 1983), distinguishing ordinary worldly occurrences (event-causal sequences) from *actions* – in which an distinctly personal-level entity, an agent, intervenes on the otherwise impersonal course of events. For Velleman, this response brings to light a real problem with the standard story but cannot be treated as the end of analysis. We need to retain, somehow, both the agent in her authorial role and the event-causal glue that attaches her to the naturalistically described world. To do this, he suggests we should aim to "locate a system of mental events and states that perform the functional role definitive of an agent" (p. 467). The aim is to find subpersonal element that is shared across actions and plays the right kind of role to justify the attribution of an action to the whole system.

The solution suggested by Velleman introduces a constitutive rationality criterion to the basic action space model. This unifies the set of representations that are to count as forming a subject's perspective on the world by limiting them to those available to those states that 'play the functional role definitive of an agent'. The definitive role of the agent, for Velleman is similar to the role, for perception, ascribed to an object by Austin (see, above §1.1b). Both object and agent serve to close a question. The question answered by agency is the question *what shall I do?* Answers to this question that are speculative or theoretical don't issue in action – they leave open the question of what an agent does or will do, partly because they are open to review in light of further considerations. When the agent acts, however, these questions are irrevocably closed. Just as perceptual presence of a thing answers the first kind of question, the agent's full-blooded intention to act in a certain way answers the second kind of question. Consequently:

> the agent's role cannot be played by any mental states or events whose behavioural influence might come up for review in practical thought at any level. And the reason why it cannot be played by anything that might undergo the process of critical review is precisely that it must be played by whatever directs that process. (Velleman 1992, p. 477)

This criterion is met by a mental state that is necessarily closed to this kind of review. The proposal Velleman makes is that this be the desire to be rational: that is, that one's actions be responsive to one's reasons. This can't be subject to internal rational critique, since any such critique would be self-defeating (to what reasons could you appeal arguing that one shouldn't be rational?). In identifying the acting subject with this desire, in which the agent "throws his

weight behind" (p. 479) his causally efficacious reasons for acting, Velleman claims to locate agent-causation in the event-causal series that constitutes action. The claim is that this is what happens when a *person* acts.

This account keeps much of the structure of the causal account set out by Davidson. The agency-constituting norms of rationality are effective inside the agent and at the moment she intends to act so that mere bodily events brought under the regime of this more complex intentional pattern. The simple belief-desire model did not capture the constitutive, agency-constituting desire that one's intention be answerable to one's reasons in general. With this in place, action-regulating norms enter into the account, and the agent as answerable to these norms enters into the causal order. Since skilled activities are governed by norms that might play a similarly constitutive role, this proposal looks promising for someone who wants to retain the event-causal framework of Davidson's model while responding to epistemic worries about agential alienation.

4) Basic Particulars in an Action Space

a) Perspectival Unity: synchronic and diachronic

Velleman' emphasis on the unifying role of rationality for agency is part of a tradition of thought about action that stretches back to Kant's famous statement of a similar criterion:

> It must be possible for the 'I think' to accompany all of my representations; for otherwise something would be represented in me which could not be thought at all, and that is equivalent to saying that the representation would be impossible, or at least nothing to me (1998, B131-132)

This observation is the starting point for Christine Korsgaard's (1989, 2009) naturalising development of a related Kantian theme: the autonomy of the agent in practical rationality. She points out that that this too involves a form of structural self-unification:

> When you deliberate, when you determine your own causality, it is as if there is something over and above all of your incentives, something which is *you*, and which chooses which incentive to act on. So when you determine your own causality, you must operate as a whole, as something over and above your parts, when you do so (Korsgaard 2009, p. 72)

The point here is that as a condition for its possibility, action selection at the personal level must be *as if it* were the self-determination of a rational (i.e., reason-guided) agent. Less than this wouldn't be intentional action, but the causal output of "a series, a mere heap of unrelated

impulses" (p. 76). For creatures capable of representation-guided locomotion and (even more so) for rational actors, this isn't how things are. Since they are appropriate targets for psychological interpretation, we can know in advance that such creatures enjoy *pragmatic* unity as a personal-level or animal-level feature of their psychological organisation, however this is subpersonally realised.

The need to determine your own causality might be thought to introduce an action selection criterion similar to the one that generate the unity of the basic action space. You need a unified perspective at any one moment to make all the acts you could perform available to you. This means need to pull yourself together in order to act, bringing a unified set of mental capacities to bear on a unified set of representations. In this case it is *as if* some kind of unified entity stood behind your action, because the separate contents of your perceptions, beliefs, and desires are brought together in determining what you actually do. The content of your perspective, in combination with your standing psychological states give you reasons for acting – some of which are causally efficacious in realised actions.

Autonomy is, for Korsgaard, something that attaches to the agent over the course of her existence – it is not a property of she exhibits at a time in order to act, but revealed in the structure of her agency as she exercises it for the whole time that she is an agent. The difference between these two can be captures in a distinction between *diaphonic* and *synchronic* agency. A collection of psychological states meets synchronic unity conditions when the states are appropriately related (to each other, to a subject) at a time – they are diachronically unified when their unifying relations hold over time. For Korsgaard, agency is unified in both senses, but synchronic unity is a consequence of diachronic unity.

This aspect of Korsgaard's view is mirrored in the interpretationist methodology set out in chapter (§2). Recall that, for interpretationists, the content (or scope) that characterises environment-directed mental states like states of perceptual consciousness depends on the subject in that state meeting conditions of interpretability (I-conditions). Subjects meet I-conditions in virtue of their behaviour, speech and, for subpersonal interpretationists like Dennett and Hurley, details of their internal cognitive organisation apt for interpretation in terms borrowed from the explanation of these personal-level, whole-agent activities. Interpretability coincides with meeting certain kinds of unity conditions – at least to a certain degree of approximation and for most of the time. For practically rational creatures, this means approximately conforming to rationally normative standards of holism such that contents attributed to the practically rational agent's mental states all have a bearing on each other and

on the activities the agent undertakes. For non- or partially rational agents, the standards may be different, or the necessary level of approximation may be lower (see Hurley 2003b). There are different accounts of what this normative rational unity, tracked in folk psychology and at least tacitly relied on by cognitive scientific explanations of behaviour amounts to; there are also various accounts of the metaphysics of psychological subjects and their states that are supposed to realise this kind of formal unity (for useful recent discussion, see Dewhurst 2021). The interpretationist approach allows us to avoid commitments on the details of these disputes while focusing on the factor that all views have in common: that there is some conception of psychological unity that ordinary interpersonal interpretation captures, and this has to do with the way that ascriptions of content to a subject's various mental states constrain each other. However the details are worked out, the fact that we can understand departures from normative unity against the background of conformity to these standards over time is instructive. The normative standards set bounds outside which we lose track of what it is to be attributing contentual states to a subject or agent, and these are established over time. This suggests that at least the minimum standards of agential unity apply to diachronically unified agents – with attributions of synchronically unified perceptions and intentions making sense only against the background of diachronic unity.

Since unity is diachronic on this view, it is what I've been calling a structural feature of perceptual and agential states (at least as far as its diachronic aspect goes). It pertains not to the content of a state at a time but to the way multiple distinct states relate to each other across time. If diachronic unity is explanatorily prior to synchronic unity, this will also mean that the latter is a structural feature: the unity of experience at a time is not part of the content of a single overall state but is a feature of how distinct states and events with their own temporal profiles contribute at any moment to a subject's unified agency.

Contrast this with the thought that synchronic unity comes first – that the diachronic unity of a perspective is built up out of a series of synchronically unified perspectives related to each other over time. This is the kind of perspectival unity offered by the basic action space model. On this model, the unity of some set of perceptual representations is determined by their joint availability to action selection mechanisms that determine moment-by-moment how an agent's temporally extended policies get implemented. So far, the claim is that internal unity of perspective is born out in the causal interaction between an agent and her world. An internal event identical with mental state onset sets off a causal chain which has as notable waypoints a) the movement of the agent's body b) the event(s) for the sake of which the action is performed. The agent is efficacious in the world through the bodily movement after which the

rest is “up to nature”. It is through these waypoints in the causal story that the action is described; we describe it intrinsically at (a) and through its causes at (b). The teleological order (a so that b) is realised internally by represented connections between the perceived scene, the desired end, and the means to attain it, and then actualised in what actually happens (a then b). This makes plausible that experience must be unified at any particular time. The unified action of the whole body must take account of what’s given in experience. If there were two or more disunified perspectives, then action would be disunified – and we can reject this for ecological reasons. Absent its role in explaining action, and finding it nowhere in introspection (indeed – this possibility seems incoherent), we have no reason to posit the further, non-unified perspective(s).

We can bring out a problem with this picture by considering the following toy-example. Imagine there were another simultaneous but non-synchronically unified perspective (call this S_1-and the first S_2). Imagine that S_1 shares content with S_2 from one moment, t_1, to the next, t_2. But while S_1 is action guiding at t_1, S_2 is action guiding at t_2. It looks like the synchronic unity of experience at t_2 depends on a different perspective than it did at t_1. But then there is nothing stopping us from introducing a third perspective, S_3, which gets into the driving seat at t_3, and so on indefinitely. We need to individuate perspectives in terms of something objective, or else we have no means of distinguishing between one and any arbitrary number of cohering perspectives inhabiting a single body. Recall that the answer proposed by the basic action space model is that intentions to perform basic actions provide this criterion. Imagine a case where four perspectives take turns to determine the movements of a single body. In this table below the action guided perspective is bolded:

S_1(P)	S_1(Q)	S_1(P)	S_1(P)
S_2(Q)	**S_2(Q)**	S_2(Q)	S_2(Q)
S_3(Q)	S_3(P)	**S_3(P)**	S_3(Q)
S_4(Q)	S_4(P)	S_4(Q)	**S_4(P)**
t_1	t_2	t_3	t_4

Table 3.1

Now we might think that it makes sense to talk of a new synchronically unified perspective emerging (S_1 at t_1, S_2 at t_2, etc.). This perspective (call it S_x) shares content with S_1, but unlike S_1 it is always causally efficacious and appropriately in touch with its environment via agency

(ignore worries about the temporal lag between environmental conditions, internal representations, and appropriate bodily responses). Does this composite perspective enjoy diachronic unity, while S_1 lacks it? This seems wrong. By stipulation, S_1-S_4 respectively stand in relations of diachronic unity with themselves. At any moment S_x, is unified with 'its own' earlier states, rather than those of other perspectives that it is accidentally coherent with. The point isn't that it's plausible to think of experience like this. Rather, it's that understanding agency according to the standard story's account cannot give us agential criteria for diachronic unity of perceptual experience, since we couldn't distinguish between multiple discrete perspectives taking turns to drive behaviour (as long as at they each had the right content *at the moment they were in control*) from a single ongoing perspective.

What's tempting at this point is to deny the first step. We shouldn't introduce S_2 because all of the content of its ongoing experience except at t_2 is epiphenomenal; it has no relation to how things are in the world and no role in directing appropriate responses. But say we get rid of all the epiphenomenal content. Since we haven't specified the connections between S_1 at different moments, it looks like S_x (cleared of the same putative connections between its separate experiential streams) is just as good at accounting for synchronic unity in behaviour. While basic action selection requires a unified experience, perceptual guidance of ongoing activity as conceived by the basic action model is compatible with a non-diachronically unified subject with merely accidentally coherent perceptual states.

Now this needn't be a problem for the standard story itself, or even to the use it is put to in accounting for synchronic perspectival unity. As I've already noted, diachronic and synchronic unity are often treated separately in the philosophical literature, so it might not surprise us that different criteria are needed for establishing them (indeed, one attempt to explain them together has already been rejected above). We might think, having established synchronic unity, the next sensible step is to work out how and when a series of synchronically unified perceptual moments come together into a single temporally-unfolding perspective. This would give us independent reasons for allowing S_1 to be diachronically unified and vindicate our intuition that S_x is not. Or we might be sceptical about diachronic unity. Maybe the argument outlined above shows that no such criteria could be given; or that, given naturalistic constraints on our theorising, none is actually met. Sceptical worries aside, there is nothing so far to suggest that a synchronically unified perspective on the world could not be constitutively linked to capacities for action, where these capacities issue in event-causal interventions on the environment.

b) Perspectival unity and basic particulars

Nonetheless, I think we should resist this picture if we can. Intuitively, there is a difference between the kind of unity of S_x and that enjoyed by the subject of experience. Partly it is this intuitive difference that talk of alienation aims to capture: S_x is, by stipulation, not attributable to any temporally continuous subject. Such a subject is therefore unable to be present in her perceptions and actions in the same way we, unalienated perceivers and agents, take ourselves to be. So, how might we resist it? One way of doing this is by falling back on introspective criteria for unity. If the requirements for action do not provide us with grounds for ruling out diachronic disunity, we might think that we can appeal to the contents of experience themselves. But this involves us in the problems alluded to in the previous chapter (§2.5c) – introspection is a relation that subjects have to content that's already given, but if the contents that we get on the table by the basic action selection criteria fail to be unified, how would introspection provide this?

A second approach would be to understand the unity of a perspective in terms of the unity of the objects that a subject has a perspective on. Strawson (1966) famously attempts to provide an argument like this, developing an argument he finds in Kant's discussion of idealism to claim that the unity of experience is provides grounds for the objectivity of perceptual content. As Hurley notes (1998, ch.. 3) this argument can be turned around. Instead of arguing from unity to objectivity, we might take the objective existence of perceptual objects for granted, leveraging this to explain the possibility of perspectival unity. Since objects are what guarantee the objectivity of experience, they allow us to make sense of subjects' ability to trace "one subjective route […] through an objective world" (1966, p. 107).

Strawson's proposals about perceptual objectivity start from a consideration of the conditions under which objective perceptually grounded statements about the environment can be made and understood. According to Strawson, these presuppose a shared conceptual structure that allows subjects to distinguish between and reidentify discrete particulars and attribute properties to these. A mapping between this scheme, the psychological lives of subjects and their shared environment is the basis for the objectivity of perception since content attributed to a subject's experience depends on the possibility of her perceptual reports being understood – and *this* depends on speaker and hearer being situated together in a world of particular objects and their properties, as characterised by the scheme.

The basic shape of Strawson's objectivity argument is as follows: first outline the conceptual scheme that allows for the identification of basic particulars. Describe the conditions necessary for perceptual reports made within this scheme to map onto environmental objects; and, on the assumption that these are sometimes met, now claim that the conditions for objective report

give the conditions for objective perception. Since this is an argument from the objectivity of perceptually grounded reference to the objectivity of perceptual content, it will be helpful to get clearer on what objectivity amounts to in these cases. Philosophers sometimes talk about the 'objective purport' of perceptual content, for example (see McDowell 2009). This is the presentation of things *as* objective: intuitively, a hallucination that seems to its subject to present the world as it is shares this contentual feature with an ordinary perception. In contrast, at least some philosophers have thought that ordinary perceptions lose their objective purport if their subjects systematically suspend judgment about any of their perceptual contents. While objective purport and objectivity in Strawson's sense are related, the target notion is one on which objective perception and objective reference are tied up with the existence of their objects. Minimally, this means that 'objective' perceptual states will be subject to two necessary conditions:

> *Objective Content*: If a perceptual state, Ψ, of a subject, S, has objective content <o is F> , then there is an object, o, in S's environment.
>
> *Objective Reference:* If a perceptual state, Ψ, of a subject, S, enables objective reference to an environmental object, o, then o is an object in S's environment and S can identify o in virtue of being in Ψ.

One way for a state to have objective content is for it to be veridical (add the condition 'and o is F') but a state can be objective without meeting this condition – one obvious way for this to be the case is if it meets conditions for objective reference. Both objectively contentful states and those that support objective reference are incompatible with the non-existence of their objects. But why think that there are any states like this? According to Strawson, a descriptive project in metaphysics entitles us to the object-side components of this picture: the particular things and their properties over which o and F range. So, the question concerns how these come to be connected with the content and possible referential uses of perceptual states. The key move in Strawson's argument is to suggest that the possibility of objective reference and content can be made sense of by showing how each depends on the other.

To start, we can look to *Objective Reference* for the beginnings of a deeper justification for *Objective Content.* The idea is that Ψ serves as means by which a subject can identify the object in the world that can be part of the content assigned to Ψ. On this picture, being in Ψ-states is part of what it is for a perceiver to have the ability to identify an object. Identification, on Strawson's early presentation of this argument is a public and linguistic act. By referring to, for example, *that* tree, a speaker can get a hearer to understand which of the range of possible

referents for the term ‘tree’ she means. Use of ‘that’ in this way can depend on a predetermined linguistic context (Strawson 1959), but the most basic use relies on the extra-linguistic context of the utterance itself. The paradigm case is pointing. A speaker refers to ‘that tree’ and if asked to specify can pick one out by pointing to it.[68]

What’s normally required for this to be possible? Both speaker and hearer must be able to distinguish some object from the overall scene and recognise it as having some property(s) appropriate to membership of the category <tree>. *Objective reference* says that perception provide these conditions; *objective content* says that the properties or features that mediate intersubjective reference fixing are ones that the objects in question really have.

Chapter (§1) has already set out reasons that we might want there to be this kind of connection between the content of perception and the possibility of perceptually grounded reference. Understood in abstraction from the possibility of reference, or granted content prior to these capacities, it is hard to understand why the content of perception should be objective. If it isn’t, then idealist alternatives threaten the possibility of ascribing particular content to people’s experiences. Strawson (1959) dramatizes this idealist threat in a discussion of the “theoretical anxiety” that the putative contents of perception as expressed in a perceptual report might fail to distinguish between a description of the objects in a perceiver’s environment and a ‘duplicate’ environment, in which a set of distinct objects with exactly the same perceptible properties existed. The case is parallel to the worry about duplicate perspectives illustrated by table (3.1) above. In both cases, the shape of the solution is determined by the possibility of objective reference:

> To experience this theoretical anxiety is to overlook the fact that we, the speakers, users of the dating and placing systems, have our own place in the system and know that place; that we ourselves therefore, and our own immediate environment provide a point of reference which individuates the network and hence helps individuate particulars located in the network. (Strawson 1959, p. 30)

What Strawson emphasises here is that both interpreters and subjects (at least as self-interpreters) already have the resources they need to secure a connection between their introspective descriptions and an objective world – this is what the moderate realism inherent to Strawson’s ‘descriptive metaphysics’ entails. These resources are not, however, to be found

[68] The theme of demonstrative identification is carried through in Evans and developed, with a greater emphasis on sensorimotor dynamics of this kind of *act* of demonstrative identification, in Hurley (1998)

in the subjectively available content of a perceptual episode alone. If the content of experience was given in introspection,

c) Anscombe on diachronic agency

In the next part of the thesis, I'll pick up on the role of objects in this story. Here, I want to focus on the role of the agent, someone who acts through a particular body situated in a particular environment. Take a paradigm exercise of the capacity to refer: the agent's pointing to an object. What can we say about this if we interpret it on the model of the standard story? In pointing to the object, one thing the agent does is refer to it *objectively*. What makes this act an objective pointing is that it is part of a communicative gesture with objective content. And what makes this content objective is the wider context of the act. An act of pointing, described as such, then violates the strict bodily constraint. Equivalently: redescribing it in accordance with the constraint requires us to isolate the bodily movements on which this description is focused as a stage on an event-causal chain that originates in a speaker's intention and terminates in her hearer's comprehension. If phenomena identified on this causal story are to be independently individuated, then the account of *objective content* just outlined will be unavailable. If we thought that the only way for communicative acts to have objective content was that we could find some component within this kind system that 'played the role of' indexing objective contents, we might experience a similar theoretical anxiety about the possibility of communication. The strict bodily constraint truncates the causal processes appealed to on this story, but then we don't think there is a strict bodily constraint on communicative actions, like pointing – if we want to talk about pointing as a basic action, we change the subject. As Strawson notes, however, we already have a solution to hand. I want to suggest that this is what we find in the interpretationist methodology set out in the previous chapter.

First, let's take stock. Anscombe denies that the events that cause a movement of the body must be identified with the psychological items that rationalise that same movement as an intentional action (§2.2b). This is because: (a) psychological causes of action are a not a unified class; any of these might be the cause of a given bodily movement, but only some are themselves apt to rationalise it; (b) rationalising psychological items are similarly disunified, and not all of them can enter into proximal causal relations with bodily actions; (c) considered as such, bodily movements cannot have *rationalising* psychological causes; even if the cause of a bodily movement happens to be a 'primary reason' in Davidson's terms, it is only situated against the background of the agent's psychological history and in its environmental setting that the bodily movement is constituted as an action.

(c) suggests a modification to the Anscombe-Davidson principle as it's been presented so far. On the Davidsonian interpretation, redescriptions of basic actions in terms of their effects (or more distal causes) do not tell us anything metaphysically relevant about the core phenomenon. It's in virtue of a certain kind of proximal cause that the bodily movement is an intentional action, independently of this wider context. Once this action is on the scene, wider action descriptions collapse into it. On the Anscombean view, movements *in context* constitute intentional actions. There is no longer a privileged description that picks out a core, action-constituting event.

This makes intentional agency essentially contextual. There is no boundary relative to which it becomes an intrinsic feature of these events that they are intentional actions:

> We do not add anything attaching to the action *at the time it is done* by describing it as intentional. To call it an intentional action is to assign it to the class of intentional actions and so to indicate that we should consider the question 'Why?' relevant to it in the sense that I have described (§ 19, my emphasis)

Assigning an event or action to this class means fitting it into an explanatory pattern that includes the agent's psychological history and past behaviour, including their participation in the practice of reflecting on, asking for, and providing reasons for actions. Certain patterns in this complex causal structure connect agents and worldly event in the right kind of way that (given these practices of explanation) those events are recognised as part of what the agent does. Anscombe and Davidson agree that understanding the agent to be situated in this context is what allows observers to recognise her behaviour as potentially subject to rationalising explanation. What Anscombe claims and Davidson denies is that actually being so situated is constitutively (rather than merely instrumentally) part of what an agent needs to be capable of intentional action.

Sensorimotor enactivists are motivated by a similar thought. Recall Hurley's discussion of the perspective constituted by the interdependence of perception and action:

> The contents of perceptual experience and the contents of intentional action both depend on a structure of causal flows that constitutes a complex dynamic feedback system […] This provides a two level account of the interdependence of perception and action that is characteristic of having a perspective (1998, p. 3)

It's an important part of the two level interdependence view that personal-level phenomena cannot be understood independently of their subpersonal implementation. Though, generally the personal level needn't be isomorphically explained by mechanisms at the subpersonal level,

and in particular sensory input and motor output to the brain do not correspond to perception and action (where these are conceived of as personal-level boundaries of the mind):

> [T]he interdependence of perception and action that perspective involves can be explained in terms of their codependence on a subpersonal dynamic singularity. This subpersonal aspect of unity does not support sharp boundaries either between mind and world or between perception and action (ibid.)

Put into the terms of the current argument: once we reject the strict bodily constrain at the personal level, we lose our motivation to find contents answering to find the contents of experience in neurocomputational mechanisms.

5) Conclusion

This chapter has set out a case for bringing together two sets of arguments about agential and subjective unity:

> A neo-Kantian theme in Hurley's discussion of the perspectival unity of perceptual experience, which seeks to analyse both diachronic and synchronic unity of experiences in terms of the ongoing agential unity of the subject of those experiences.
>
> An Anscombean theme in recent action theory, which denies that the 'bodily component' of an action can be non-arbitrarily specified apart from the agent whose action it is: an agent to whom the question 'why?' could in principle be addressed.

Both of these tendencies push against assumptions that make isomorphic subpersonal explanation tempting. Both propose structural or formal unities (of the intention and its manifestation in action, of the subject across and at a time) and interdependencies (between what an agent does and the context she is situated in, between the perceptual affordances and effects of action) as an alternative to this approach. It proposes these as the personal-level targets for subpersonal explanation, which precedes not by finding counterparts for them at the lower level of analysis but by guiding and constraining subpersonal content attribution in light of this understanding of the phenomena. Rather than asking how the subpersonal (synchronically) constitutes this personal-level structure, this approach asks how each component of the mechanisms appealed to contribute to the achievement of these unities and dependencies over time and in context.

A positive answer to the questions about the unity of perceptual experience from which this chapter started goes has now become available: what pulls together the different strands of

perceptual experience – the different *experiences* out of which episodes of overall perceptual experience is composed – is their shared, agency-mediated contact with the environmental objects. Because the agent must act in this environment, and because action depends on a unified perspective on it, we are right to attribute unity to experiences where they are available to guide and inform an agent's intentional actions. Working out this account in more detail is the job of the next part of this thesis, which will argue that the subject of unified perspectival experience has a mode of acquaintance with environmental objects, and that this determines its content and phenomenal character.

Part II

The Sensorimotor Acquaintance View

4

Objects of Perception and Environmental Objects

0) Introduction

The next two chapters set out and defend the sensorimotor acquaintance view, the positive account of perspectival experience that the rest of the thesis has been building towards. First, let's review how we got here. Chapter (§1) gave the desiderata that the sensorimotor acquaintance view aims to meet. Perspectival experience situates subjects in their environments, it is what allows us to trace a subjective route through a world of objective things (Strawson 1966) and so realises a basic kind of perspectival (Hurley 1998) acquaintance with that world. But it constitutively depends on internal, neurocomputational mechanisms that are causally insulated from that environment. A good account of its phenomenal character should respect both of these facts. It should also clarify what it is to perceptually encounter a world. In part, this means saying what falls within the 'scope' of perspectival experience.

Some contemporary naïve realists think this can be done by appealing to a higher-level relation of acquaintance between subjects and environmental particulars. Objects of acquaintance are supposed to play a constitutive role that's different from the one played by neurocomputational mechanisms, providing the extra ingredient needed for answering this scope question. While the debate has been advanced by this way of accommodating the cognitive scientific literature, less attention has been devoted to articulating just how the distinct constitutive roles of personal-level acquaintance and subpersonal mechanisms fit together. Chapter (§2) attempted to fill this gap by defending a two-level interdependence model of subpersonal explanation, along the lines suggested by Hurley (1998, 2008). Subpersonal explanation maps a psychological vocabulary characterised by interpretive norms of holism and environmental appropriateness onto neurocomputational mechanisms; subpersonal content ascription depends on personal-level interpretation, but interpretations are constrained by mechanisms. Explanatory isomorphisms occur where subpersonal explanations maximally constrain interpretation, so that parts of the psychological capacity or phenomenon can be uniquely mapped onto the components of the neurocomputational mechanism. Fully generalised, this picture yields a 'flat'

or ‘one-level’ cognitive science. Many local controversies in philosophy of mind and cognitive science turn on the availability of explanatory isomorphisms of this kind, and we can usefully think of the debate around naïve realism as one of these cases.

In chapter (§3), I claimed that explanatory isomorphisms are unavailable for perspectival experience. Having a perspective involves practical awareness of the interdependence of perception and action over time. Perception and action situate subjects among environmental particulars and being so situated amounts to having a unified perspective on the environment. Since this interdependence is causally mediated by things in the environment, proposed two-level isomorphic models of either capacity ignore its dependence on the other – no overall isomorphism between subjective perspectives and neurocomputational mechanisms is available. This suggests that some environmental particulars have a special role to play in these experiences’ phenomenal character. By mediating sensorimotor interactions on which perception and action depend, these coordinate the subpersonal processes underlying perspectival experience in a way that explains their proposed status as objects of acquaintance. This and the following chapter assess this possibility and gives a qualified positive verdict. As we’ll see, to adopt the sensorimotor acquaintance view we must abandon some characteristic epistemic commitments that have traditionally supported naïve realism.

The main job of this chapter is to give an outline of the sensorimotor acquaintance view. This provides the basic shape of the proposal, which will be elaborated and defended in chapter (§5). Since this view gives a central role to the objects of perspectival experience, it is important to give an account of what these are and how they come to be related to the subject of experience. This chapter focuses on these questions, defending an account of the environmental object of experience based in the cognitive scientific notion of a perceptual object. In section (§4.1), I restate the semantic motivations for the sensorimotor acquaintance view, highlighting a similarity between these and those involved in naïve realist models of acquaintance. I contrast this with the package of semantic, epistemic and phenomenological commitments behind standard naïve realism. I motivate an approach that takes perceptual objects as they are characterised in perceptual psychology and related fields to be candidate objects of acquaintance. In (§4.2) I introduce the notion of a perceptual object, defending a mild regimentation of our ordinary talk about the objects of perception in the face of an objection from Austin. In (§4.3), I review contemporary empirical work on capacities for reference and object tracking, connecting this to the parts of the philosophical picture already on the table. In particular, I discus the ‘perceptual object file’ construct, around which much of contemporary cognitive scientific work on object perception is organised.

1) Appearances and Perceptual Objects

a) Object-appearance properties and phenomenal character

I've already suggested that naïve realists face difficulties when they try to pull apart the semantic, epistemic and phenomenological motivations for their characteristic claims (§1.1-1.2). If representational sematic or epistemological theories could account for the sematic or epistemic particularity of perception, plausibly this would undercut the introspective motivation for the view as well. Since philosophers have been led in various directions by their introspection-based intuitions,[69] it looks like we at least need a principled reason to think that introspective support for naïve realism, as opposed to other views, is reliable. An excellent reason for thinking *this* would be if perceptual experience's introspective transparency were a source of knowledge – knowledge of the intrinsic properties of perceived objects. This is why I've read Martin as sharing the epistemic motivations I've associated with the Oxford Realist tradition, and in particular with Austin (Austin & Warnock 1962). But if perceiving an object in the right conditions entails knowledge of its intrinsic properties, it seems that this knowledge in turn entails a capacity to refer to the object and its properties. So, it looks like all three factors come together in the naïve realist picture, even if what gets the project going is an attempt to capture the phenomenal facts, as revealed to introspection.

Whether or not it can withstand the various objections from cognitive science reviewed in previous chapters, this package falls apart when the introspectable phenomenal properties at its centre are assigned to anything other than the environmental objects of perception. This is the crucial problem with Beck's (2019) 'neurocomputational naïve realism', which understands appearance properties as intrinsic to the neurocomputational states mediating subjects' awareness of objects. If appearances are separable from external objects, then attention to them, even under conditions of introspective transparency, cannot be a direct route to knowledge of environmental objects. There has to be a further story about what takes the subject from appearances given in the acquaintance relation to properties intrinsic to the object. But this is just what the appeal to acquaintance was supposed to provide.

[69] Martin (2002, p. 376) notes that there's something vaguely paradoxical about this: "It is common to think that how things appear to one is something obvious to oneself – or at least that it should be obvious as long as one is suitably attentive to the question. So, one might ask, how can there be sustained debate about what is obvious? […] Although it is puzzling that there can be any debate about appearances, such debate clearly does exist." Put in terms introduced in chapter (§1), what this indicates is that commitment to *transparency*$_{ins}$ gives rise to an additional explanatory burdens. Naïve realists need to account for the existence of the philosophical debate and find a way to privilege the evidence provided by 'naïve' introspection.

As what's required is a mode of relation between perceiver and object, a natural strategy for naïve realists is to appeal to the actual spatiotemporal position of the object relative to the subject. The object of acquaintance appears as it does, partly because of how it is physically situated relative to the perceiver. Part of the story about why anything or sounds as it does is because it is seen (heard) from *here* and over (or at least out) *there*. This is partly what Campbell has in mind when he characterises the perceptual standpoint (see, e.g., Campbell 2002, pp. 19-20; Campbell & Cassam 2014, pp 27-28); and we've now seen that this approach has precedents in the work of the later Oxford Realists (Strawson 1959, 1966; Evans 1982). But, in this case, why not think that it's the neurocomputational representation of these properties that does the explanatory work? Neither Strawson nor Evans were naïve realists, and more recently Dickie (2010, 2015) has articulated an argument along similar lines for a kind of perceptual acquaintance that does not give perceived objects a constitutive role in phenomenal character. Why not think that objects appear as they do because of how their spatial and temporal properties are *represented* in experience, along with their other perceptible features (cf. Burge 2005, 2011)?

I think that our best answer to these questions follows Hurley's (1998) suggestion that the role assigned by this tradition of acquaintance theorists to objects' first-order spatial and temporal properties is really played patterns of interaction mediated by sensorimotor agency. These are a function of objects' locations in time and space and these representation-guided mediating processes, and so cannot be reduced in either direction. Since the relevant sensorimotor patters are partially determined by subpersonal representations of objects' spatial and temporal location, this proposal resists the internalisation of appearances while accommodating the empirical results on which Burge's (2005, 2011) and Pautz's (2013, 2023) critiques and Beck's (2019) response are based.

At the centre of the sensorimotor acquaintance view's alternative proposal, then, is the claim that phenomenal properties are relational properties of the objects of perceptual acquaintance. These are properties they have because of their participation in sensorimotor interaction of this kind.[70] The next few paragraphs give an outline of this picture, highlighting basic claims and providing introductory exposition.

[70] Since it assigns phenomenal properties only to the objects of the acquaintance relation, the view is in some ways closer than Beck's to the full naïve realist package just outlined. As he notes (Beck 2019, ft. 24) when naïve realists talk about 'appearance properties' they're more often referring to a subset of the perceptible properties of perceived objects (see, e.g., Genone 2014).

First two schematic claims about appearance properties themselves. One is about appearance property instances:

> *Object-appearance property*: An object, o, instances an appearance property, A_P, iff for a subject S there is a subject-object interaction Sϕo, where Sϕo appropriately yields the interpretation <o appears P to S>

Since appearance properties are real properties of objects in the environment, the notion of appropriateness that fills this schema has to come from a discourse that aims to describe the behaviour of perceiving subjects in an objective world. As we've seen above, both folk psychology and cognitive science aim to do this; and I've suggested that the latter succeeds in meeting this aim by drawing the former's explanatory resources. This means that norms of appropriateness in attributing appearances to subjects (and so appearance properties to objects) are determined by the practices personal level explanation in folk- and scientific psychology. Chapter (§2) sketched these, but it did not aim to codify these norms – on one hand, because they are part of the experimental practice of a developing branch of science and, on the other, because they emerge in the perennial, everyday task of interpersonal interpretation. Rather, what's key for the present argument is that these two sets of interpretive norms exist and come together in the kind of cross-level explanations required to answer questions about the structure of perspectival experience.

Granted that there are some such standards for appropriates, we have the basis for a second claim – about how we can characterise types of objects by their appearance properties:

> *Object-appearance type*: An appearance property P characterises an object type O_P iff there is a range of possible circumstances C_1, …, C_n which appropriately yield interpretations of the form <o appears P to S>

These circumstances (C_1, …, C_n) need not be actual and will in every case be paired with possible interpretations of interactions of the relevant kind. Interactions are *at the focus of* the wider circumstances relevant to their interpretation, in the sense introduced in (§2.2). For now, we can think of these interactions in intuitive personal-level terms: as attendings, noticings, or observings of objects in environmental contexts, against the background of their subject's psychological history and in service of their current activities (§2.1-2.2). Instancing an appearance property is having an appearance. But belonging to a type characterised by an appearance property needn't be having an appearance or an appearance property, since some or all of the relevant values for C might be determined by non-actual circumstances. While, for

example, describing something as 'red-in-ordinary-conditions' invokes possible circumstances relative to red (or red-looking) objects, this schema makes explicit that subjects and their modes of interaction with objects in the circumstances are relevant to determining appearance types.

Appearance properties are complex relational properties of objects. Different given values for S, ϕ and o will determine different appearance properties and so different corresponding object-appearance types. While differences in some of the intrinsic properties of a given object will have systematic effects on appearance properties when the values of S and ϕ are held fixed, intrinsic properties do not determine appearance properties in general. Appearance properties are the objects' contribution to the phenomenal character of episodes of perspectival experience. Whatever else needs to be said about these episodes' phenomenal character, appearance properties instanced by objects participating in them will exert a systematic effect. This is the aspect of phenomenal character that is captured by interpretations of the form <o appears P to S>. Attaching properties to interpretations like this is the start of a subpersonal explanation, since they characterise a class of interactions that can be a target for content ascriptive mechanistic decomposition.

The second pair of claims are about the phenomenal character of episodes of perspectival experience:

> *Appearance contribution:* An object's appearance properties constitute an aspect of the phenomenal character of an experience that provides a subjective perspective on that object.
>
> *Non-object exclusion:* No aspect of an experience that does not provide a subjective perspective on an object has a phenomenal character constituted by object appearance properties.

The idea is that objects make a unique contribution to the phenomenal character of experiences that bring them into a subject's perspective. Perception and agency lock onto objects in a way that gives a characteristic structure to episodes of perspectival experience. Combined with our background principles linking phenomenology to personal-level content (§2.1) and personal-level content to interpretation (§2.2), this generates the constitutive claims expressed in *appearance contribution* and *non-object exclusion.* By specifying an object's appearance properties, you single out some aspect of the phenomenal character of an experience of that object: this is (an aspect of) the object's appearance.

It should now be clear that 'appearance properties', as I'm using the term, introduces a semi-technical notion that abstracts away from several of the senses that 'property' and 'appearance' have both in ordinary and more specialised contexts. Since it's going to be put to so much work, this might worry us. Especially coming after a discussion of the Oxford Realist tradition, we might here remember Austin's admonition against attempts by philosophers to improve on our "stock of common words" (Austin & Warnock 1964, p. 4). As Austin points out, where philosophers introduce their own terms to solve philosophical problems, they can fall prey to a kind of circularity: characterising sense-data as introspectively available objects of perception and then appealing to introspection of perception to support existence claims about these objects, for example. Fortunately, we needn't rely on philosophical fiat to introduce appearance properties. A similar property and property-characterised object type is already in use in psychology and related disciplines, where classes of 'perceptual objects' are routinely discussed in terms of the abilities of perceivers to detect and respond to environmental particulars. So, while we might characterise a perceptual object as any environmental particular that a subject can perceive, we look to empirical work on perceptual capacities to specify what these are.

The rest of this chapter refines the basic claims set out above, connecting them to a subpersonal story about object-perception. But first, I'll try to head off a few more sceptical objections to the shape of the programme just laid out. This will take the form of some brief clarifications of what the target notions of appearance and appearance properties *aren't*.

So understood, the notion of an object-appearance property *does not*:

- Give an analysis of 'objective phenomenal properties' (e.g., redness, hardness, etc. *attributed to objects*) in terms of their dispositions to produce experiences of a certain corresponding character. Complex causal dispositions of objects link them to effects that can be described at the subpersonal level, but these causal features are part of what *constitutes* an the phenomenal character of experience.
- Give an analysis of 'subjective phenomenal properties' (e.g., redness, hardness, etc. *attributed to 'qualia'*) in terms of their standard environment causal source. Appearance properties are non-intrinsic objective properties of perceived objects.
- Give a way to exhaustively describe perceptual phenomenology. The view does not say that appearance properties are all that there is to the character of perspectival experience, and it does not deny that there could be aspects of phenomenal character that are 'ineffable' or inaccessible to interpretation. Since the proposal is only about the appearance properties that show up in perspectival experience, adopting it is

compatible with having commitments either way about other aspects of phenomenal character, including putatively ineffable ones.

- Give a way to pick out a 'natural kind' of environmental object. Though some natural kinds of object will be associated with some appearance properties, appearance property instances are context sensitive in a way that is likely to make them unsuitable for characterising a natural kind type.

While arguments could be constructed that link appearance properties to kinds and properties along these lines, they don't play a role in my case for the sensorimotor acquaintance view. As this case is pressed bellow, considerations will emerge that tell against at least some versions of these possible arguments (though I won't set them out systematically). For now, we only need to bear in mind that nothing below presupposes the controversial or dubious conclusions of any such argument. Note also that since the above claims don't generate a requirement that appearance properties are instanced by any independently identifiable types of environmental object, we have scope to deal with various kinds of unusual appearances in non-standard conditions. This is helpful in responding to common philosophers' intuitions about 'hallucinations.' When the scenarios that generate these intuitions are described in enough detail, they turn out either to be impossible or to involve interactions with highly ecologically unusual objects (see §1.1).

Nonetheless, it's worth emphasising that appearance properties (and so perceptual object types) are determined by objective features of the interactions between perceiving subjects and their environments. And the appearance properties that we talk about and use appeal to in explaining and understanding aspects of these subject-environment dynamics come from discourses that seek describe these dynamics. An accurate description of an object as having an appearance property depends on the availability of an accurate description of a subject as responding to some objective property (possibly an abstract or higher-order property) that the object has. As we've seen, this might be thought to lead to a kind of circularity: doesn't every object have the 'higher-order property' of eliciting responses in subjects that are appropriate to the kind of object it is in some given context of interpretation? What saves the proposal is that there are prior, external constraints on what kinds of discourse can set the standards of appropriate interpretation. The biconditionals given above are not meant to introduce or define a new class of properties, but rather to bring out some common features of the way that folk-psychological and cognitive scientific discussion of subjects' responses to the environment deal with appearances. Sciences have one way of earning their place in a system of interpreting perceivers and agents, folk practices have another, but both must avoid the kind of empty or

circular explanations just alluded to. Since the standards of appropriateness that give us the appearance properties we actually use in such explanations exclude such circularity inducing properties, they are also excluded from the philosophical account set out below.

In a more positive vein, I'll now state some natural consequences of the way that appearance properties have been characterised so far:

- There may be some more-or-less idiosyncratic appearance properties, so that the property is instanced only in relation to one perceiver, or only to a narrow range of perceivers.
- There may be some universal appearance properties, so that being a perceiver entails being able to have experiences constituted by properties of this type.
- Some appearance property types will have a determinable-determinate structure that crosses over equivalence classes constructed by fixing values for S. So, e.g., an object that appears P to two subjects might appear P_1 to one and P_2 to the other, where P_1 and P_2 determiners of P.

As we'll see, these features of the view offer a satisfying response to some traditional worries about conflicting appearances and introspective reflection on phenomena of perceptual constancy.

Moreover, since perceptual object types are introduced by reference to the abilities of real-life embodied subjects and these abilities are explained by mechanistic explanation, we can turn to these lower-level explanations to explain why certain kinds of perceiver perceive certain kinds of object. Some explanations of this kind are pushed beyond the subpersonal level, for example by consideration of the properties of an animal's sensory array, scientists can say what kinds of physical properties it can respond to. Since the intrinsic properties partly contribute to appearance properties, those that depend on properties that an animal is insensitive to cannot be perceived by that animal.

This consideration helps in responding to metaphysically based worries about explanations that appeal to relational (as opposed to intrinsic) properties of objects. The kinds of relational properties to which the sensorimotor acquaintance view appeals need be no more (or less) metaphysically troublesome than relational properties introduced in these and other 'special sciences' (cf. Hurley 1998, pp. 34-35). Granted them, these relational properties are what give us object-appearance properties; and appearance properties give us object-appearance types. Ultimately, in any perceiver-object interaction, what determines that a perceptual object that

falls into one such type is the sensorimotor interaction itself. So, the fit between perceivers and perceptual objects is a consequence of the fact that object-appearance properties ontologically depend on these interactions. Only once it is possible for perceivers to bring their sensorimotor capacities to bear on an object that the object can have the relational properties that get the sensorimotor acquaintance view's proposal going. Of course, the object apart from the interaction has intrinsic properties, and these make a difference to what interactions it can participate in. Fundamentally, then, it might well be that the relational properties are fixed by the intrinsic properties of objects, perceiving animals and their environments. But the fundamental level of description at which these intrinsic properties would be ascribed is neither the personal or subpersonal at issue for psychology and cognitive science.

Nonetheless, and despite the other qualifications sketched above, an account of the phenomenal contribution of perceptual objects advanced on suitably developed versions of these claims cannot be adopted without cost. Having laid my cards on the table, I'll now review the stakes. If what accounts for the phenomenal character of experience are properties that objects have only in virtue of participating in sensorimotor acquaintance relations, the epistemic component of the naïve realist package comes into question. While, in chapter (§1), I suggested that the phenomenological distinction that Beck tries to draw between a subject's being appeared to in some way and an object's appearing that way to a subject is unmotivated (given the way he characterises appearance properties), naïve realists need something like this distinction. On one version of the naïve realist picture, knowledge fills the gap. Being appeared to in some way cannot be a source of the knowledge that something is that way unless appearances are transparent (i.e., *both* $\text{transparent}_{\text{Ins}}$ and $\text{transparent}_{\text{Ext}}$) to objects' intrinsic properties; and so, if acquaintance *is* knowledge of these properties, then it must be transparent to them. A simple story about how this could be possible is that an object's intrinsic properties determine its appearance properties. The sensorimotor acquaintance view does not provide a story of this kind. It gives us no reason to expect a phenomenological difference between two episodes of experience of objects whose intrinsic properties differ, but which participate in sufficiently similar sensorimotor interactions with perceivers. While many have found claims of this kind intuitively plausible, naïve realists have rightly pointed out that they entail epistemic limitations on perception that are more difficult to accept (see, e.g., Campbell 2002; Gomes & French 2016).

I think we should accept these limitations.[71] This means that the sensorimotor acquaintance view must subscribe to a less ambitious vision of what acquaintance can do for the subject of experience. Though it doesn't entail or presuppose it, holding the view is compatible with thinking that, in Russell's terms, perception leaves us "completely in the dark about the true intrinsic natures of physical objects" (1946, p. 18).[72] By understanding acquaintance as a ground for reference rather than a guarantee of knowledge, we can make sense on this limitation on perceptual capacities while endorsing the claim that the contact with objects that acquaintance affords gives a characteristic phenomenological shape to episodes of perspectival experience.

2) Perceptual Objects

a) Introducing perceptual objects

To motivate this proposal and to situate it against the historical backdrop laid out in previous chapters, let's consider how something like the overlapping notions of 'perceptual objects' used by cognitive neuroscientists, psychologists and ethologists could help with resolving questions about perception returned to by successive generations of Oxford Realist philosophers. First, a characterisation. Very roughly, the category of perceptual objects includes all the things in the environment that perceivers can single out or discriminate from one another without the assistance of cognitive abilities – as in, for example, applying sortal concepts or descriptions. So, shadows, dots on a screen, sounds are plausible candidate perceptual objects, as are organisms and ordinary material things. Football teams, things that belong to the emperor, visual *Ganzfelder* and arbitrarily selected spatial or temporal parts of perceptual objects are not.[73] More controversial items to include on either list would be things like silences and absences, distinguishable spatially or temporally continuous parts of objects and components of complex sounds. In introducing the notion, I'll not assume a resolution either way for difficult cases like these.[74] I'll also restrict my discussion to perceptual objects that show up in perspectival experiences through their visible or audible properties. While some discussion of perceptual

[71] That is, we should give up on this distinction between 'being appeared to' and 'being appeared to by o' at the level of phenomenology. In characterising the epistemic situation of the perceiving subject, we might still appeal to a distinction between 'mere appearances' and genuinely knowledge-conducive encounters; but we would have to deny that these are associated with any distinctive phenomenological contrasts.

[72] The view leaves open the intuitively plausible option that perception does secure knowledge of objects intrinsic properties *given other conditions of the subject.* The claim is that phenomenology is determined by appearances and that these are, by themselves, not knowledge-conducive.

[73] Or at least for some of these, are not *under normal conditions.* If the things that belong to the emperor are all bundled together or moved as a unit or if the emperor only owns one thing, then (depending on how you resolve some nearby questions in metaphysics), they might be or become perceptual objects. Similar things might also be said for the football team, and so on.

[74] For discussion, see Green (2018)

objects treats them as internal representations of items listed above (see, e.g., Shinn-Cunningham 2017; Barwich 2019), I'll follow what I take to be the mainstream usage.[75] So, in what follows, 'perceptual objects' refers to a class of things in the environment and 'perceptual object representations' or (more specifically, in a sense to be introduced below) 'object files' refer to subpersonal representations of perceptual objects.

From very early stages of development onwards, perceptual objects get and hold attention; throughout life they are the targets of a very large proportion of everyday perceptually guided thought and action (Johnson 2011). Ecologically, this makes sense because perceptual objects largely coincide with ordinary material things and organisms. Tracking perceptual object often enough means tracking these. Getting things right about them means getting things right about the kinds of things that most urgently bear on survival and reproductive success (Pylyshyn 2007; Green 2019). Consequently, questions about subpersonal and personal-level capacities to represent and respond to perceptual objects make up the main theme of several overlapping research projects in perceptual psychology and neuroscience (Green & Quilty-Dunn 2021). Researchers characterise personal level capacities for object perception in terms of peoples' ability to respond to and selectively attend to things in the environment. Across object type and sensory modalities, perceptual objects are associated with a profile of behavioural responses that are specific to the class. I'll say more about how perceptual objects are characterised in the psychological literature below, focusing on 'object files' as a proposed subpersonal mechanism underlying these capacities. For now, we should note three features of perceptual objects: a) they are materially constituted by parts of the mind-independent external environment; b) characterising type-identity conditions for perceptual objects involves (implicitly or explicitly) appealing to the perceptual capacities of at least some kind of perceiver;[76] c) they can be the objects of *objective perceptual reference* (§3.4).

We might expect (a) and (b) to furnish material for an explanation of (c). For example, we might think that the mind-independent material constituents of perceptual objects account for their ontological independence from experiences of them; while properly spelling out the relationship between perceptual object and capacity types alluded to in (b) would make sense of why perception puts subjects in a position to refer to objects. So, (a) explains objectivity; (b) explains reference. A further natural thought is that perceptual capacities work by representing some of the properties objects have just in virtue of their material constitution. Call these the Intrinsic

[75] Mainstream, at least, in the visual literature. Green Quilty-Dunn

[76] Consequence: a statement of the constitutive conditions of being a perceptual object must make reference to the perceptual capacities of at least some possible animal.

properties of perceptual objects. Taken together, these considerations might lead you to conclude that perceptual content (at either personal or subpersonal levels, or both) is objective because it correctly attributes intrinsic properties to objects, and that this is why (a) and (b) explain objective reference.

Below, I'm going to resist this line of argument and provide an alternative way to realise the intuitive formular – (a) and (b) explain (c) – just outlined. This account drives a wedge between perceptual objectivity and knowledge of intrinsic properties. But for now, I'll just assume that perceptual objects can be directly referred to by someone who gets their intrinsic properties wrong – that objective content is not a necessary condition of objective reference. This view is defended by semantic acquaintance theorists, like Dickie (2015). A stronger claim, which Dickie rejects but the sensorimotor view entails, is that perceivers could refer while getting objects' intrinsic properties wrong *systematically.* The second claim is what will be argued for in section (§4.5), but we only need to assume the weaker claim to motivate an account of acquaintance which targets psychologists' perceptual objects.

b) Motivating case: Decoy ducks vs. duck data

Austin (Austin & Warnock 1962, pp. 78-83) criticises a distinction made by Ayer (1940) between 'existentially delusive' and 'quantitively delusive' perceptions. Roughly, the distinction tracks the way many contemporary philosophers use the terms 'hallucination' and 'illusion'(see, e.g., Siegel 2011, p. 34; Martin 2002). *Existentially* delusive perceptions are supposed to underlie experiences that are subjectively indistinguishable from episodes of object perception when there is in fact no object perceived by the subject; *qualitatively* delusive perceptions have real objects, but they attribute to these objects features or properties that the objects don't really have. In the terms introduced by the previous chapter (see §3.1), qualitatively delusive perceptions fail to satisfy conditions for objective content, but they do satisfy conditions for objective reference. Existentially delusive perceptions fail to satisfy both sets of conditions. The criticism Austin wants to make in these passages is that the contrast between just two types of 'delusion' is an arbitrary and unrealistic regimentation of the ordinary talk about perception on which it depends. Paying attention to the various distinctions speakers actually make between 'appearances' and 'reality' reveals central cases that do not fall under this rubric.

Here's Austin's discussion of what he takes to be a particularly important type of case like this:

> Suppose that I see a decoy duck[77] and take it for a real duck; in which of Ayers ways is my 'perception' to be said to be 'delusive'? Well, it just isn't clear. It might be held to be 'qualitatively' delusive, as endowing the material thing with 'qualities that it does not really possess'; for example, I might mistakenly suppose that the object I see could quack. But then again, it might be said to be existentially delusive, since the material thing it seems to present does not exist; I think there is a real duck before me, but in fact there isn't. (Austin & Warnock 1962, p. 79)

Austin uses the term 'material thing' for the purported object of this experience as of a duck. For reasons that we've already explored (see §1.2, §1.4), he's resistant to a way of talking that separates this material thing – the duck/decoy duck – from its appearance. By this point in his discussion, he's rejected one way of doing this, according to which 'material things' are one side of a pair with 'sense data' (see ibid., pp. 3-5).[78] This is Ayer's preferred expression and framing, but since Austin doesn't put the term in scare quotes – as he does with 'perception' (used as a count-noun) just before – we needn't read him as adopting it in his exposition of Ayre's view.[79]

If we take it that the sense of 'material thing' in this passage is up for grabs, we can say that part of Austin's point is that we don't have a way of characterising the thing seen that naturally describes the situation Ayer has in mind without begging the question (try, for example 'the *apparent* duck'). For Austin, the decoy duck illustrates a class of cases "in which I think I see something where something else really is" (p. 79). If we try to specify the putative object of perception in terms of what the subject thinks she sees, then it doesn't exist; if we appeal to whatever object *is* there, then the object and the appearances trivially come apart.

Even a sympathetic commentator, however, might here register a concern that Austin is overly quick to move from these observations to his conclusion that "in the case of the decoy duck [Ayer's distinction] breaks down at once" (ibid.). Using the notion of a perceptual object just introduced, we have a way to articulate this worry. The difference between cases in which someone non-delusively sees a decoy duck and cases in which they see a decoy duck and take it

---

[77] A decoy duck is a model duck, sometimes used as a lure in duck hunting. Not to be confused with a 'duck decoy' which is a kind of trap for ducks and other waterfowl.

[78] On Austin's argument against sense data, "what is spurious is not one term of the pair, but the antithesis itself" (1962, p. 4). Contrasting this approach with the contemporary mainstream rejection of sense data views we can get a sense of what's behind Martin's (2007, pp. 7-8) comment that "where philosophers have come to reject a commitment to sense-data, this rejection is the most shallow form of religious conversion, repudiating no more than a form of words" While standard representationalist positions reject sense data, they don't avoid talk of material things in the way Austin recommends.

[79] He might be adopting it – the argument still works, as far as it does, if he is. But I think running the 'ordinary language' reading in the main text makes more sense of the problem as Austin sees it.

to be a real duck is that in the second case they mistakenly attribute a property (or cluster of properties) that ducks have and decoy ducks don't *to a perceptual object*. So, this is a case of qualitative delusion, in something like Ayer's sense. We're entitled to say this because we have a way of characterising the class of perceptual objects that ducks and decoy ducks belong to. We can take it that (under certain conditions) the visual system of some relevant range of perceivers will pick up on features that ducks and decoy ducks have in common, allowing them to be discriminated from other objects (i.e., non-ducks) but not from each other.

These features – I'll call them perceptual 'cues' – are at least partly determined by the material constitution of token duck-like perceptual objects, but not always in the same way or for the same reasons. Notably, they're realised differently by ducks and decoy ducks.[80] They also depend on features of the environment in which perception occurs (it's only under certain conditions – e.g., within a certain range of relative distances between duck and viewer, given certain lighting conditions, etc.) and on aspects of the functional organisation of perceivers themselves that objects provide distinctive cues. This last heading might include relatively more or less stable properties of perceivers, from basic features of the perceptual system to contingent differences in attention (after all, if you *really pay attention*, you're less likely to mistake a decoy for a duck).

Let's call these factors taken together a 'perceptual situation' and the relationship between a situated object and the set of perceptual cues that the object presents in the situation specified by the perceptual situation cue-from-object dependency. This gives us a way of characterising perceptual object types:

> *Perceptual object type:* A token perceptual object, o, belongs to a perceptual object type O_F iff the cue-from-object dependencies of o's perceptual situation deliver cues in a range characteristic of F-s enough of the time.

What counts as *enough* of the time, is likely to vary across contexts of investigation. So, for example, a study investigating decoy duck effectiveness (see, e.g., Eaton et al. 2015) might try to determine whether the cue-from-object dependencies of two kinds of decoy made a difference to the behaviour of ducks. Objects would fall into different categories if they elicited

[80] The green colour of a mallard duck's neck, for example, is partly determined by the way specialised structures in its feathers diffract light. This is not what determines the greenness of a decoy mallard's neck, which is green because of the chemical properties of the plastic it's made of or the paint it's painted with. This leads to perceptible differences, which are even greater in UV ranges to which birds but not humans are sensitive – though this doesn't seem to matter for the effectiveness of ordinary decoys (Eaton et al. 2015).

statistically significant differences in the behaviour of duck population being studied. Implicitly, this would give evidence on the range of cues characteristic of perceived ducks: one kind of decoy eliciting more duck-appropriate behaviour is evidence that it delivers more duck-like cues. The use of decoy ducks to lure real ducks exploits the fact that they fall into an equivalence class of this kind. A good decoy duck is one that in a range of appropriate ecological conditions is a source of enough duck-like perceptual cues to elicit flocking behaviour in real ducks.

It's worth bearing in mind here that no claims are being made about the status of these categories except that they are put to genuinely explanatory use outside of philosophy and they entitle us to the distinctions Ayer is thinking of. An 'apparent duck' will be a member of an equivalence class that includes both real and decoy ducks. While the boundaries of these classes are likely to be fuzzy and context sensitive, they impose real constraints on membership that determined by actual and possible interactions between class members and a corresponding kind of perceiver. Relations like these are in the background of any interaction between a perceiver and an environmental object. In cases like Austin's duck, there will be at least one equivalence class into which the object of a 'qualitative delusion' and a corresponding good case of perception both fall. This won't be the case for 'existential delusions'. If we leave out cases in which the perceiver simply fails to respond to relevant cues at all, these options exhaust the field of possible failures of object perception.[81]

Perceptual object types characterised by cue types play a distinctive explanatory role in perceptual psychology and ethology, a role played by environmental objects that deliver perceptual or sensory cues and not played by arbitrary interventions either on cue-from-object dependencies or after sensory registration of cues. By contrast, the distinction between 'sense data' (the kind of thing that we perceive), and 'material things' (that kind we don't) is a philosophical or, to use Austin's preferred epithet, 'scholastic' distinction. It contains and tacitly expresses the philosophical theory in support of which it is appealed to. Both categorisation schemes are theory-laden, but only in the second case is the corresponding theory conceived at an *unearned* level of abstraction and distance from the phenomena it seeks to explain. . So, even if we're convinced by Austin's claim that "[t]here is no one kind of thing we perceive but many different kinds" we should note his qualifying remark that the number of such kinds is "reducible

[81] This gives us a way to identify the object and characterise the delusion in terms of properties misattributed to it. The material thing that's present to the subject in the bad ('delusive') case is a duck-like perceptual object, as in the good case. What's more, in the case where she rightly takes herself to see a duck, she's also presented with a duck-like perceptual object, the duck. Ducks are, in a sense, *paradigm* duck-like perceptual objects. But articulating what this sense is and why they are will have to come later, when more of the pieces are in place.

if at all by scientific investigation and not by philosophy" (Austin & Warnock 1962, p. 4; see also p. 127). While appeals to sense data might be philosophy in bad faith (Travis & Kalderon 2013), the notion of a perceptual object gains currency from the use it has been put to generating insights unavailable from non-empirical reflection.

But even if we can appeal to perceptual objects to ground Ayer's distinction between two ways object perception can go wrong, there's an important disanalogy between them and 'material objects' in his sense. Perceptual objects are introduced and type-individuated in terms of humans' and other animals' capacities to perceive them. While one could try to integrate Ayer's model with psychologists' discussion of perceptual objects (e.g., by talking about decoy ducks' dispositions to produce 'duck data'), this would be unmotivated since sense data don't come up in this discussion. And, as Austin notes, it would be orthogonal to Ayer's main project, which is motivated by considerations about introspective transparency (Austin & Warnock 1962, pp. 102-103). This aspect of the proposal is the target of well-known objections. But while very few contemporary philosophers or cognitive scientists appeal to entities like sense data in personal-level models of perception or experience, other aspects remain mainstream. As naïve realists complain, a lot of contemporary work on phenomenal consciousness analyses phenomenal character in terms of the intrinsic phenomenal properties of psychological states. The relation between these and environmental objects outside the 'Markov blanket' of their neurocomputational correlates, would be generically causal (cf. Hohwy 2016). We're left with a similar puzzling disconnect between content and phenomenal character and between the subject and the objects of experience. Focusing on perceptual objects, and their potentially constitutive role in perceptual phenomenology addresses both of these complaints.

3) Object Files

a) Object Files and Perceptual Objects: First Pass

Since we've introduced and characterised perceptual objects in terms of cue-from-object dependencies whose causal dynamics are partially shaped by subjects perceptual capacities, getting clearer on what perceptual objects are requires us to say more about the capacities that target them. This will allow us to set out the principle behind the lists of examples given above. Ideally, this should also give us a way of dealing with difficult cases; but we will have made progress even if we only arrive at a deeper justification for the inclusion of paradigm instances.

At a coarse grain of description, it should already be fairly clear what sort of answer we're aiming at: a visual object belongs to a type discriminated by the visual system, a token visual object will be something that reliably is picked out by sight or is just now being seen; auditory objects can

be heard, an auditory token is a potential or actual object of audition. So, perceptual objects are not just anything that perceivers or their perceptual systems pick out or respond to. Their intrinsic properties and environmental context need to contribute to generating distinctive cues for perceptual systems. This rules out hallucinatory percepts and purely endogenous phenomena like those accompanying migraine aura (see §2.3), since in such cases nothing in the environment is perceived.

It also suggests a plausible general characterisation of capacities for object-perception: object-perception is the ability to respond appropriately to cues. Standards of appropriateness are determined by computing a 'reverse function' for the relevant the cue-from-object dependencies (Kersten et al. 2004). If particular objects determine particular cues, we should expect these cues to lead back to their objects. But, while as a matter of fact they generally do, we can't take this for granted. It's a basic insight of perceptual psychology that the patterns of input that deliver cues to perceptual systems massively underdetermine the objects of perception (ibid., Fodor & Pylyshyn 1981). Since patterns of input considered by themselves undetermined their objects, an explanation of why cues are *cues for* a given object must appeal to mechanisms that selectively responding to that object. We can get an initial handle on subpersonal explanation of this target capacity by asking *what is required for the perceptual system given cues to arrive at representations of corresponding objects?*

For about the last 30 years, a significant part of mainstream psychology and neuroscience's answer to this question has appealed to perceptual 'object files' (Kahneman et al. 1992; Green & Quilty-Dunn 2021). The earliest formulations of the object file construct were in work on vision (Kahneman et al. 1992; Pylyshyn 2000, 2007) where it continues to sustain the largest empirical literature (for reviews, see Scholl 2001; Carrasco 2011; for discussion, see Green 2017; Green & Quilty-Dunn 2021). Investigation of auditory object files has had nearly as long a history (Kubovy & Van Valkenburg 2001; Griffiths & Warren 2004; Nudds 2010). More recently, discussion of multi-modal perceptual object files has been at the focus of a lot of work in cognitive science (Jordan et al. 2010; Zmigrod & Hommel 2013; and philosophy (O'Calaghan 2015; Green 2019) and a related research programme has proposed corresponding 'event files' (Hommel 2004) for action and sensorimotor integration (Lacey et al. 2009). Although I'll draw mostly on results from the more established visual and auditory research programmes, the account of perceptual objects that I end up endorsing supports natural extensions into other modalities – as well as at least a strong presupposition in favour of auditory-visual object files (and corresponding objects). Sensorimotor event files will be discussed separately in the next chapter (§5.3). But it's worth noting that the sensorimotor acquaintance view is in-principle separable from the details

of these models. For example, the view is compatible with cognitive scientists finding that distal objects and bodily activities are subpersonally represented in radically different ways.

Perceptual objects and perceptual object files are best understood as coordinated notions. Perceptual files function to pick out perceptual objects; perceptual objects are those things that object files pick out when they're working as they should. Slotting object files into the formula for perceptual object types arrived at in the last section allows us to introduce the following sufficient condition on perceptual objects:

> *Object file sufficiency*: o is a perceptual object if the good functioning of an object file is characterised by its representing objects of a type that o is a token of

Should we upgrade this into a statement of necessary and sufficient conditions? There are several reasons to think that the resulting claim would be too strong. Let's briefly consider two.

First, depending on how the relevant notion of function is understood, the types of object that object files function to represent might be restrictively and implausibly limited (e.g. to objects typical of relevant selection environments; see, e.g., Burge 2009).[82] If function explains content, this means that there's a problem with many of the contents assigned to object files in the empirical literature: either because standard cases involve pervasive misrepresentation or because contents aren't those that have standardly been attributed (Green 2018). Strict teleofunctionalists about subpersonal representational content might be happy to accept these consequences – or might reformulate the notion of function to avoid them (see, e.g., Milikan 2004, p. 113-127) – but the present argument shouldn't turn on how such disagreements are resolved. As we've seen above (§2.1, §2.4), our interpretationist framework licences a way of talking about subpersonal content that allows us to be noncommittal here. If we want to do this, we should be open to the possibility of there being perceptual objects that are explicitly represented by object files even though there is no object file type whose good functioning requires its doing so. That is, even if we think that object files are crucially implicated in all human object-perception, we needn't tie subpersonal content to function in the way that would require us to defend an object file necessity claim.[83]

Secondly, pending a fuller characterisation of perceptual object files, it is unclear how widely they are realised by non-human perceivers that we would pretheoretically want to say are 'in

[82] For a classic discussion of this notion of function, see Millikan (1984)

[83] We still should want to charaterise object in terms of function though – otherwise we can't get below the personal interpretive level, since we won't be able to state what an object is in terms any more precise than the ordinary language way we individuate objects of perception.

touch with' the same set of perceptual objects as us (or are in touch at least with an overlapping set). Assume, in line with claims sometimes made within the literature, that object files of some kind are a pervasive feature of the perceptual systems of all mammals (Brefczynski-Lewis & Lewis 2017) or even all vertebrates (Clark 2004). This would still leave us with a staggering number of creatures that don't in some important sense share a perceptual world with humans at all. Even if we're hardnosed enough to accept this, related considerations ought to give us pause. At the very least, we should want to make space for the possibility that cognitively sophisticated subjects with perceptual systems radically different to humans could refer to a substantially overlapping set of perceptual objects.[84]

So, while the stronger claim would give a more direct route from the notion of an object file to a model of acquaintance along naturalistic lines (and might straightaway force judgement on some of the hard cases listed above; cf. Dickie 2010), I think we should reject it. And, as we'll see below (§4.3a), the considerations raised against the strong claim will also give us reasons to accept a permissive characterisation of the perceptual objects that the sufficiency claim gets on the table (cf. Green 2019). With this now in hand, we can start to get a sense of what these might be by considering what object files are invoked to explain and how they are supposed to do this.

b) **Object Files in Psychology and Cognitive Neuroscience**

In their recent philosophical review of the object file literature, Green and Quilty-Dunn (2021, p. 667) give the following general characterisation:

> [O]bject files are complex mental particulars consisting of indexes and short-term memory stores in which features of the indexed object are encoded. The index is what 'links' the information in the store to a particular object. Such information is accurate or inaccurate by virtue of being about a particular object, and the index determines the object against which it is assessed for accuracy.

As some philosophical commentators have noted, the distinct components proposed by these models (indexing and feature-representing mechanisms) can be assigned the two "distinct logical functions" described in Strawson's (1950, 1959) theory of reference:

[84] Asking just how different these subjects might be and correspondingly just how far we should expect perceptual objects to overlap raises a several difficult questions that I won't address here. Minimally, though, we shouldn't want to stipulate that a creature without an object-file system couldn't perceive any of the objects an ordinary human can. Talk about the 'inscrutability of reference' – causal account going to be helpful here.

> One is proto-predicative: the capacity to sense red (or any other feature) both here and there. The other is proto-referential: the capacity to identify the red region as one that is also stippled. This is an informative identity, grasped sub-personally. According to these models, property binding—successful feature integration—consists in sub-personal grasp of such identities (Clark 2004, p. 450)

We'll see below that these features have been proposed as the subpersonal basis for a Strawsonian argument linking the possibility of object representations and the capacity for non-descriptive reference to mind-independent spatiotemporal particulars (Campbell 2002; Dickie 2010, 2015, ch. 4). That is, they've been suggested as the subpersonal conditions for one half of the unity-objectivity argument discussed in the previous chapter (§3.5). Looking ahead to the next chapter: in (§5.3) I'll argue that combined with the non-isomorphic sensorimotor approach outlined over this thesis, they contribute to a version of the whole picture.

Notably, researchers of both vison and audition emphasise the dual role that object files are supposed to play in singling out and assigning properties to environmental particulars by combining more basic features of the sensory input. Here's a representative example from the literature on vision:

> An object file is a mid-level visual representation that "sticks" to a moving object over time on the basis of spatiotemporal properties, and stores (and updates) information about that object's surface properties. (Gao & Scholl 2010, p. 83)

Here Gao and Scholl distinguish between two types of 'properties': the visual system's response to 'spatiotemporal properties' are what make it possible to index an object to an object file, while its 'surface properties' are what enter the file's memory store. These are supposed to be a set of objective characteristics that are 'directly available' in conscious vision – colour, shape, and motion (Kahneman & Treisman 1984; Pylyshyn 2000). By linking these properties to environmental particulars, object files "help us to construct our conscious perception of objects and how they behave – e.g., telling us 'which went where'" (Gao & Scholl 2010, p. 2010). These properties survive periods of occlusion and changes in the transitory properties which enable the file to pick out and track it. For convivence, let's call the former 'properties' of objects and the latter 'features' (cf. Pylyshyn 2007).

Pioneering models of the visual object file system, like those developed by Anne Triesman and colleagues (Treisman & Gelade 1980; Treisman 1988) characterise a 'linear' processing hierarchy: feature binding comes first, which enables indexing of later occurring features to an

established perceptual unit, with the evolution of features indexed to a single object allowed for property attribution. But while these models have clear intuitive appeal for vision – as well considerable empirical predictive validity (Hochstein, & Ahissar 2002; Kristjánsson & Egeth 2020) – they are faced with obvious difficulties in other modalities. For example, auditory cues provide *mostly*[85] temporal information, and for this reason putative auditory object features are largely dynamic rather than spatial. Even in the visual domain, simple linear processing models have come under pressure from more recent empirical results (for discussion, see Humphreys 2016).

Further controversies revolve around the role of selective attention in these models. Attention is widely agreed to play an important role in the functions I've associated with object files, but there are disputes about to what extent it is necessary for feature binding, and corresponding disputes about how far feature binding is a separable capacity, exercised prior to object representation (see Burge 2022, ch. 5). While some complicated features are thought to be represented in processing areas before those responsible for object files, there are currently unresolved questions about the relation between these earlier processes (call them 'feature integration') and 'feature binding' in the above sense.

To avoid stepping too far out onto any of these empirical limbs, I'll limit myself to the following basic picture: pre-attentive feature integration occurs at retinotopically organised early visual processing areas and multiple features are indexed to a single perceived entity in an object file. These files can be said to 'construct' perceptual experiences of objects by providing targets for visual attention and by carrying out feature-to-property computations (however these functions are related).

Similarly, auditory object files – sometimes called 'auditory event files' (see, e.g., Zmigrod,& Hommel 2008)[86] – are supposed to single out particulars and assign properties. Here's a typical

[85] The qualification here relates to the so-far imprecise notion of a 'cue' which *could* be understood as the spatially distributed effects of a sound-wave on a hearer. In an ordinary human this might be thought to yield a two-dimensional spatial element (corresponding to the two ear canals, cochlea, etc.). However, on all standard models, early auditory feature processing does not explicitly represent spatial features or exploit purely spatial information (Bizley & Cohen 2016) – though this picture is complicated in the usual ways by the possibility of top-down effects on early perceptual processing (Tóth et al. 2016 Zmigrod & Hommel 2008). These issues notwithstanding, 'sound location' is typically understood as a property of auditory objects.

[86] More than in the visual literature, researchers into hearing tend to use the term 'object' to refer to the subpersonal representation or the subjective percept of a sound – and sometimes to both – rather than to the environmental entity responsible for auditory stimuli. So, for example, in a review of neuroscientific work on 'auditory objects', Bizley and Cohen (2013, p. 693) characterises them as "the computational result of the auditory system's ability to detect, extract, segregate and group the spectrotemporal regularities in the acoustic environment into stable perceptual units". Nonetheless, for the sake of clarity, I'll follow the moderate linguistic reform suggested by Green and Quilty-Dunn (2021)

characterisation of how features of a sound sequence give rise to the identification and attribution of objective properties to a distinct auditory object:

> The key to this process is the formation of a representation which captures the regularities common to a coherent sequence of sounds; a 'model' of a putative sound source […] However, in addition to encoding a regularity, this representation is predictive of the sounds that the source is likely to emit and hence can underpin the formation of an identifiable perceptual unit (object) as well as its separation from other units. (Winkler et al. 2009, p. 534)

According to this picture, which extends Bregman's (1990) influential two-stage model of the subpersonal mechanisms of conscious hearing, pre-attentive audition combines temporal features of sound sequences into potential 'perceptual units' or 'proto-objects.' These can be discriminated from the wider 'auditory scene' and associated with identifying sources (see also Shammer et al. 2011). Note that these sources needn't be identified with perceptual objects in other modalities, although the circumstances where this does happen are the focus of ongoing research (Zmigrod & Hommel 2013). A 'source' can be a placeholder indexing multiple sounds or sound sequences to a single perceptual object (Winkler et al. 2009; cf. Hommel 2004).

Bringing these parts of the empirical picture together, we can now more precisely characterise three main functions of object files:

a) Assigning multiple distinct features to a single object at a time. The object file brings together multiple feature representations realised at early 'pre-attentive' processing stages and assigns them to a common referent.
b) Assigning multiple similar and distinct features to a single object over time. The object file keeps track of the evolution of features assigned to its referential content, across changes in initially assigned features and motion in egocentric space.
c) Ascribing stable properties to a single object over time. The object file assigns and keeps track of objective properties and location in absolute space that survive changes in represented features and relative spatial location.

Taken together, this constellation of functions is what I've been referring to as an object file's capacity to 'pick out' or 'represent' a perceptual object. But it will also be convenient to have separate labels for each. I'll call (a) 'feature binding', (b) 'object tracking', and (c) 'property

---

and reserve the term 'object' for the (materially) mind-independent target of object files (cf. Clark 2004, p. 463). Since neuroscientific 'auditory objects' represent particulars and assign properties, I'll call these 'auditory object files.

attribution'. While these labels are drawn from the empirical literature, there is some variation in how different researchers use these and similar terms. Without trying to adjudicate these differences, this list stipulates the terms' meanings for current purposes.

By now there's a large evidence base favouring the object files framework over alternative theories of mid-level perceptual processing in both auditory and visual domains. While I won't attempt to reproduce the empirical case for the framework, some salient points of this experimental literature are highlighted in *this note*.[87] The picture generated by this converging evidence is that a large part of perception is organised around segregating and selectively processing discreet units, which largely correspond to common sense objects, events and sounds in ordinary environments. This goes some way towards justifying an approach that draws from it. To the extent that any theorising about the role of perception in people's psychological lives has to take on certain empirical commitments, we should want to accommodate at least the more central and uncontroversial elements of the currently most productive and unifying framework (cf. Dickie 2015, pp. 115-116).

4) **Conclusion**

This chapter has outlined and defended a view of the object of perspectival experience. These are elements of the structure of a perceiver's environment non-arbitrarily picked out by mechanisms of object perception. They correspond to the Strawsonian notion of a basic particular, in that they are ontologically independent of the subject but are constitutive of the conceptual scheme on which her ability to think about, describe and act within the world builds. Capacities for object perception are grounded in the good functioning of these mechanisms. But if this good functioning is understood according to the non-isomorphic framework set out in previous chapters, the objects they pick out have a special explanatory role to play – in the

[87] Evidence from studies on multiple visual object tracking (MOT) and auditory scene analysis suggests that perceptual objects are associated with specific processing limits. Subjects can easily report several features or properties of a small number of objects, but performance drastically reduces after they're required to target further objects above this processing limit – this performance inhibition is dissociated from the effect of increasing target features while keeping object number constant (Awh et al. 2007) Reaction speeds are improved when congruent features are placed on 'reidentified' visual objects. Attention 'shifts' more easily within a single visual object than across objects (Behrmann et al. 1999). Work with fMRI consistently correlates visual (Xu & Chun 2006; Fang et al. 2008) and auditory (Arnott et al. 2009) networks and processing areas with dissociable processing of binding, tracking and property attribution. Auditory objects have also been associated with distinct electrophysiological signatures in EEG and MEG studies (Sussman et al. 1999; Ding & Simon 2012; Tóth et al. 2016) and similar signal components have been identified for visual feature binding (Gray & Singer 1989; Tallon-Baudry & Bertrand 1998; Kaiser et al. 2004).

activation and sustaining of object files and, as I'll argue below, in the structure, content and phenomenal character of episodes of perspectival experience.

5

Accomplishment and Acquaintance

0) Introduction

This chapter elaborates and defends the sensorimotor acquaintance view, filling in the details of the story began in the previous chapter (§4.1). The sensorimotor acquaintance view comes in two parts. The first is a claim about the subpersonal basis of perceptual reference. Perceivers' capacities to speak, think and act in a way that is responsive to particular objects in their environments depends on the integration of internal representational and sensorimotor abilities, in the first instance, to distinguish objects from their surroundings and track their movements. Distinguishing and tracking objects supports perspectival experience when representation guided perceiver-object interactions take on the structure captured by personal-level singular thought and intention attribution. Achieving this structure is becoming perceptually acquainted with the perceived object. The second part connects acquaintance and phenomenal character. By determining how sensorimotor interactions play out, perceptual objects shape the episodes of perspectival experience they participate in. The appearance properties of the objects of acquaintance are what bring about these systemic effects and so constitute objects' contribution phenomenal character of corresponding perspectival experiences. Adopting this account involves rejecting alternative proposals that locate the phenomenal contribution of the object in its intrinsic properties or in internal representations of these properties. In so doing, the sensorimotor acquaintance view makes available a principled, non-introspectively grounded answer to the scope question raised in chapter (§1): the contents of perspectival experience are fixed by the relational properties of perceptual objects that shape subject-object interactions.

In (§5.1), I review Imogen Dickie's (2010; 2015) attempts to fit the traditional objects of acquaintance into this subpersonal story. Dickie defends a semantic acquaintance model that might be thought to shore up the argument for phenomenal acquaintance. However, her development of the proposal relies on the kind of interlevel explanatory isomorphisms that screen off environmental objects from the subpersonal constitution of perceptual states. In (§5.2), I distinguish between this aspect of the proposal from another part of the view, which is more promising for phenomenal acquaintance theorists. This is the accomplishment model of acquaintance. In (§5.3) I argue that developing the accomplishment model on the basis of sensorimotor capacities gives a constitutive role to the environmental object of perception,

allowing their perceptible properties to shape the phenomenal character of perceptual episodes.

1) Acquaintance as Epistemic Skill: Dickie's Semantic Acquaintance Theory

a) Perception, reference and semantics: overview

There's an obvious connection between the functions and contents ascribed to object files in the cognitive scientific literature and the personal-level capacities for direct perceptual reference at the focus of contemporary discussion of perception and perceptual acquaintance. The content that researchers across this field assign to object files is *singular* – object file representations are accurate or inaccurate about a particular thing in the environment, if at all. Perceptual object files index or 'stick to' persisting perceptual 'units' and attribute properties to them. They are closely associated with conscious perception and are given a key role in explanations of perceptual constancy (Cohen 2015; Burge 2022, ch. 3), as well as fundamental personal-level capacities to refer to objects in language and thought, formulate reports of their perceived appearances and objective properties, and perform environmentally appropriate actions.

As a result, some acquaintance theorists think that this work can help resolve some of the long-standing puzzles that the personal-level discussion revolves around. The thought is that 'proto-predicative' and 'proto-referential' format of object files (cf. Clark 2004) offers the basis of a subpersonal explanation of whole-agent capacities for perceptually grounded reference. This strategy is at the centre of Dickie's defence of 'acquaintance-theoretic' accounts of perceptually grounded reference against 'descriptivist' rivals (Dickie 2010, 2015, chs. 2-4; cf. Campbell 2002). While Dickie's contribution is primarily a thesis about perceptually grounded belief, it involves substantive claims about perception and experience as well as the relation between these and their neurocomputational realising mechanisms. By taking a detour into semantic acquaintance guided by her sophisticated application of a pattern of argument closely related to what I've been calling isomorphic subpersonal explanation, this section aims to bring out some more detail of the target phenomenon and the methodological approach.

The descriptivism/acquaintance theory debate is at the focus of seminal and ongoing work in philosophy of language, parallel to the debates in philosophy of mind and cognitive science that this thesis has so far been engaging with (see §1.1). Descriptivists about perceptual reference claim that a necessary condition on being able to refer to an object on the basis of perception is being able to grasp, understand or use a description that ascribes a perceptible property to that

object. Acquaintance theorists deny this. Typically, they think that grasp of the description is to be explained at least in part by a prior relation of acquaintance with the object itself.

As with phenomenal acquaintance, the terms of this debate can be traced back to Russell (1910). Russell was a semantic acquaintance theorist about reference to sense data and a descriptivist about reference to things in the environment. Similarly, Moore's account of sense data was paired with a 'partialist' (Clark 2006; cf. Snowden 2015) theory of perceptual reference. According to this, reference to whole objects is mediated by direct reference to the parts of those objects that are in view (or are heard, felt, etc.), with the descriptive content of beliefs or expectations taking up the slack. Discussing the empiricist tradition that informed the sense data theories of Russell, Moore and Ayer, Dickie (2015, p. 116) diagnoses the "baroque miseries of their various attempts to show how perception enables thought" as symptoms of inherited empiricist assumptions about what the subject takes from perception. But if we reject the assumptions underlying this approach, we lose our grounds for corresponding restrictions on the possible objects of acquaintance. Consequently, we lose one reason for trying to understand reference to these objects on the model of concept-mediated description.

A more positive reason to include everyday environmental objects in the scope of acquaintance comes from reflection on the relationship between perceptual phenomenology and perceptual demonstratives. First, consider this description from the phenomenologist Maurice Merleau-Ponty:

> If I am walking on a beach toward a boat that has run aground, and if the funnel or the mast merges with the forest that borders the dune, then there will be a moment in which these details suddenly reunite with the boat and become welded to it. As I approached, I did not perceive the resemblances or the proximities that were about to reunite with the superstructure of the ship in an unbroken picture. […] Afterwards I recognised as justifications for the change the resemblance and contiguity of what I call "stimuli", that is the determinate phenomena obtained from up close and with which I compose the "true" world. "How did I not see that these pieces of wood were part of the boat?" (pp. 17-18)

This first-person description highlights many of the features of perception that the sensorimotor account is tailored to accommodate. Subpersonally, the phenomenal unity of the boat described by Merleau-Ponty can be explained by feature integration in an object file corresponding to a perceptual object. But this integration is facilitated by bodily action – the subject's intentional movements relative to the boat. These movements shape the subjective episode, which moves

from a less to a more secure acquaintance with the thing perceived, suggesting a reciprocal role for elements at the two levels of explanation in generating this perspectival acquaintance. Despite the heightened register of Merleau-Ponty's description, the experience is familiar: moving to bring something more clearly into view.

I'll return to these resonances with the semsorimotor view below; but in order to get the second element of the relationship – perceptual reference – on the table, it will be helpful to modify Merleau-Ponty's presentation:

> *Maurice and Bertie:* Imagine that the beach stroller described in Merleau-Ponty's vignette has brought someone along with him. Call this pair Maurice and Bertie. 'Look how tall that is!' says Maurice, pointing towards the mast of the boat, and is surprised when Bertie replies that it doesn't look particularly tall to him. Bertie, who up until this point had been attending neither to the ship, its mast or the treeline behind it is seeing things as Maurice had some moments before. Maurice has undergone, but Bertie hasn't, the 'reorganisation of appearances' that pulls the boat's mast out of the screen of trees behind it and fuses the two together.

What Maurice means is that the *ship* is tall. Bertie, to whom the boat without its mast doesn't look tall, doesn't agree. This might be for one of two reasons: a) Bertie understands that Maurice means the ship but doesn't think the ship is tall because it doesn't look tall. B) Bertie thinks Maurice is pointing to a tree, but since the mast Maurice is pointing to blends in with the treeline behind it, it doesn't stand out as particularly tall. It looks more-or-less as tall as all the trees around it (in fact, since *it* is the mast of the ship, it is somewhat shorter and somewhat closer).

It's worth noting that while there are conceivable circumstances that would yield each of the two interpretations of Bertie at (a) and (b), scenario (a) is more realistic.[88] This casts Bertie in the same position Maurice has just been in. Correspondingly, the (a) interpretation of Bertie's state of mind should strike us as more intuitive. Imagining oneself into Bertie's point of view, it's easier to picture Maurice seemingly pointing to a not particularly tall ship than pointing to an arbitrarily selected tree behind the ship. If the putative tree had been tall enough to stand out, meanwhile, it's easier to imagine Maurice seeming to indicate it but correspondingly harder to

[88] There's empirical evidence that this is the verdict of common-sense 'folk' psychology and that this verdict is supported by the subpersonal mechanisms of perception (Moll & Meltzoff 2011). I think all of this contributes to the case for the interpretationist methodology I'm adopting (see §4.4; cf. Hurley 2008).

imagine mistaking it for a tree in the first place. This, at least, is what I take to be the most natural pre-theoretical picture of the situation to be.[89]

We might think a natural descriptive content for Maurice's statement will look something like this:

M: 'That *ship* is tall'

Correspondingly, the natural descriptive contents for what Bertie understood Maurice to have said in scenario (a) would look like this:

$B_{(a)}$: 'That *ship* is tall'

$B_{(a)}$ expresses a claim that Bertie takes Maurice to have made and a belief that he takes Maurice to have in scenario (a) (assuming that Bertie takes the claim to be sincere, and no other unusual conditions obtain). Since he doesn't share the belief, he denies the claim – the content shared by *B(a)* and *M* is the subject matter of Bertie and Maurice's disagreement. The less-natural-seeming content of Bertie's understanding in scenario (b), meanwhile, would look like this:

$B_{(b)}$: 'That *tree* is tall'

In this case, disagreement is only apparent, because there's no shared subject matter under discussion. Bertie takes Maurice to be talking about something other than what he's really talking about.

In these sentences, 'ship' and 'tree' are sortal nouns. They group particulars into types by their possession of properties and features. Often enough, some of these are perceptible. So, sticking with our example, but relativising hearers and hearer meanings to scenarios (a) and (b), perceptual demonstrative fixing reference *might* work like this:

> *Simple descriptivist diagnosis:* Bertie_a, understands Maurice's use of *M* because he understands it as having the same meaning as a sentence of the type $B_{(a)}$, and takes $B_{(a)}$ to refer to a particular object, o, because o has enough of the perceptible feature of ships and o is located roughly where Maurice is pointing; Bertie_b misunderstands Maurice's use of *M* because he understands a sentence like $B_{(b)}$, and takes $B_{(b)}$ to refer to a particular object, o', because o' has enough of the perceptible features of trees and is located roughly where Maurice is pointing.

[89] A classic, theory-based worry here is that in cases like these the reference of 'that' is 'inscrutable' (Quine 1960)

An initial problem with the *simple descriptivist diagnosis* is that since o and o' are both located in the same rough location vis-à-vis Maurice and Bertie, we don't yet seem to have an explanation for *why* Bertie_a understands Maurice and Bertie_b misunderstands him. But, of course, descriptivists *do* have an explanation for this: Bertie_a uses his grasp of $B_{(a)}$ to determine *what* Maurice is pointing at – and it's likewise with Bertie_b and $B_{(b)}$.

This way of fixing the *simple descriptivist diagnosis* generates the more fundamental problem that acquaintance theorists find in descriptivist theories in general. Sentences like $B_{(a)}$ and $B_{(b)}$ contain the demonstrative 'that', which belongs to a family of words (including, e.g., 'this', 'there', 'them') that for present purposes we can call 'deictic terms.'[90] The role of dietic terms is to indicate some object in a context-relative way. Maurice can talk to Bertie about *that* ship, because Maurice and Bertie are together appropriately located in the same environment. Meanwhile, *this sentence* can refer to *that* (same) ship because of the linguistic context generated by, narrowly, the last few paragraphs and, more widely, the quoted text. If Bertie grasps Maurice's reference descriptively, there will be a basic feature in common between these two cases: the context that determines the reference for the deictic term will contain an individuating description of the thing referred to.

Let's reformulate the descriptive content of sentences *M*, $B_{(a)}$ and $B_{(b)}$ along these lines, removing the dietic term 'that', so that *M* and $B_{(a)}$, for example, share the content:

M_d: 'There is something that is a tall ship'

Read 'there' in M_d as part of an existential quantifier over things that are tall ships rather than as a demonstrative indicating a location ('there is a tall ship' not '*there* is a tall ship'). Now we have a further problem. *M* refers to a particular ship, but if we want to explain the singular content of *M* in this way, M_d won't do. This is because M_d could be used to correctly describe a scene with arbitrarily many tall ships in it. If Bertie had understood *M* to mean M_d then he might not have disagreed with Maurice, while taking o to be a non-tall ship, as long as something in sight *was* a tall ship. Both Maurice's original remark and its translation as *M* are only true if something is tall where Maurice is pointing. For descriptivists, the content of Bertie's reference-fixing description is supposed to guide his thought to the thing Maurice is pointing to, and the

[90] More strictly, linguists and philosophers of language talk of deictic *uses* of words – of which, the spatial use of demonstratives is only one kind (Wolter 2009). But, as we'll see, the term has had an uptake among cognitive scientists and philosophers in something like this sense, used of subpersonal mechanisms of perception (e.g., 'diectic codes'; Goodale 1997).

content of M_d doesn't do that.[91] What we need is a *unique, locating* description, so that anything that has the features and property that the sentence ascribes to something *has to be that thing*. We need a description more like this:

M_{ud}: 'The ship Maurice is pointing at is tall'

But now, two more problems: a) Descriptivists about deictic reference can't help themselves to individuating names like 'Maurice' without giving a descriptivist account of reference fixing for these; b) the most plausible replacements for 'Maurice' (e.g. 'the thing *he's* pointing at') introduce further indexical terms.

There are moves that descriptivists can make in response to these points, but we can leave the discussion here: (a) and (b) label two more domains in which the 'baroque miseries' evoked by Dickie play out. The upshot of all this is that descriptivists have one kind of job and acquaintance theorists have another. Descriptivists have to specify the content of Bertie's reference fixing description is in a way that doesn't fall foul of these objections. Acquaintance theorist have to determine what the non-descriptive conditions for singular reference are. A natural place to look in cases like *Maurice and Bertie* would be the conditions of joint attention to perceptual objects. It's natural for the same reason that scenario (a) has a more intuitive feel to it: on looking at the scene, with no other objects to attract his attention, we would expect Bertie's gaze to be drawn by the boat, rather than by some arbitrarily selected element of its immediate background.

b) Acquaintance and perceptually grounded reference

One can approach questions about perceptual acquaintance with an agenda informed by prior commitments in this debate. For example, the claim that perceptual acquaintance is a non-representational psychological relation can be derived from semantic acquaintance theorists' characteristic negative commitment: if personal-level perceptual states have representational content ascribing their object with individuating perceptible properties, then a main pillar of the descriptivist position is established.[92] But acquaintance theorists can leave this pillar standing if they provide reasons to think that subpersonal perceptual representations are unsuitable for playing the role that descriptivists require of them. This might be because a) they don't

[91] In all this, I'm assuming that the colocation of Bertie and Maurice in the same spatial frame of reference is implicit – there's an implicit 'here' on the end of sentences like *M* and M_d. This is not an unproblematic solution (cf. Strawson 1959; Campbell 2002), and descriptivists will have to say something about objections along these lines too. But, if they can answer the worry about distal reference it's plausible that they could draw on their solution to the one worry to address the other.

[92] Not the whole story though – still need to say how the content individuates its object, how subject's grasp of content enables linguistic reference.

individuate the environmental objects of perception; b) they don't individuate particulars *by description*; c) they don't determine the personal-level contents on which a subject's individuating description would have to draw. The strategy for defending the sensorimotor acquaintance view denies (a) and accepts (b) and (c). But it says that (c) is more fundamental: we should accept (b) partly *because* we accept (c).

Here's another proposal. Take the fundamental claim to be (a). The contents of subpersonal object files don't refer to their objects, rather they provide resources for genuinely individuating reference that takes place at the level of the whole subject. Feature binding and tracking provide one kind of resource, securing an information channel from an environmental object to the subject. Property attribution provides another, registering patterns of reliable correspondence between features and object properties. But it requires the conscious attention of a perceiving subject to bring these subpersonal resources together and realise that subject's perceptual acquaintance with an environmental object. Acquaintance, a form of perceptual knowledge of particulars, is a phenomenon of the personal level.

This is roughly what Dickie argues we should take from cognitive scientific work on object files. On this view, being the target of subpersonal binding, tracking and property attribution is sufficient for being a perceptual object and that being a perceptual object is necessary for being an object of acquaintance. The extra, personal-level ingredient is provided by a knowledge condition on objects of acquaintance:

> *Perceptual Knowledge Necessity*: o is an object of perceptual acquaintance only if perception is a means of knowing some intrinsic property (or properties) of o

Given these background commitments, *Object File Sufficiency* and *Perceptual Knowledge Necessity* jointly entail that objects of perceptual acquaintance are all and only those things whose intrinsic properties can be known through perception. If perceptual objects are things in the environment that object files pick out and attribute properties to, this means that objects of acquaintance are those for which property attribution is a ground for knowledge.

This proposal supports the claim that perception does not refer by description without relying on a specific, controversial account of the personal/subpersonal distinction like that provided by chapter (§2). It therefore supports a defence of perceptual acquaintance that isn't hostage to any potentially controversial independent argument for (c). Instead, since (b) follows trivially from (a) and (c) follows from (b), we secure all of these claims at once. By itself, this won't take us all the to full-spectrum naïve realism. Dickie does not endorse an acquaintance theoretic

account of perceptual phenomenal character, and in several places distinguishes her proposal from perceptual naïve realism (see Dickie 2015, pp. 119-120). But it disarms the divide-and-conquer approach suggested by Siegel (2019) (see §1.1). So, we might think that naïve realists can adopt Dickie's proposal and supplement it with some further claims – with these together yielding the full naïve realist picture.

There are reasons to think this won't work. Dickie's proposal rests on a model of perception as a source of subpersonal informational input to personal-level cognition. But this picture makes sense only within an explanatory framework that isomorphically maps the sensory and motor periphery of subpersonal processing to the outer boundaries of perception and action. As I've argued above, naturalistic relational accounts of phenomenal character must reject this picture. This needn't be a problem for Dickie since she isn't committed either to relational account of perceptual appearances or to methodological naturalism. But, while I won't offer arguments that could decide the underlying naturalism anti-naturalism dispute (cf. Dickie 2015, pp. 12-14), I will claim that a 'naturalised' version of the proposal secures the semantic and epistemic payoffs of acquaintance in a way that blocks off the route to a relational account of perceptual phenomenology.[93]

First, let's review the features and commitments of Dickie's view:

> *Content uptake*: The target of object files' indexing and tracking functions are ordinary environmental objects. Beliefs formed by uptake from object files refer to the objects indexed by those files.
>
> *Non-descriptivism*: Perceptual object files perform their indexing and tracking functions in a way that is conceptually prior to property attribution. Property attribution depends on indexing and tracking, so that, e.g., indexing an object, o, does not require the system to represent o as having some property F, but representing o as F requires indexing o.
>
> *Epistemic homing*: Perceptual reference to o is justification-supporting convergence of information gathering resources on o. Beliefs about o are justified when they are

[93] It might be worth noting that anti-naturalism is at least as controversial as the non-isomorphic style of subpersonal explanation that supports the sensorimotor acquaintance view. Dickie's *full* proposal denies what the sensorimotor view denies – i) that subpersonal perceptual content individuates by description and ii) that personal and subpersonal levels coincide. But the rejection of (ii) is based on this framing anti-naturalistic assumption and as a result is compatible with interlevel isomorphism (see §2.1). My claim is that a naturalised version would require it.

supported by ways of gathering information that are reliable with respect to the properties of o. The availability of these resources to S enables S to refer to o.

Non-naturalism: All of the above claims can in-principle be defended by appealing to special proprietary features, abilities or properties of psychological subjects or perceptual systems. For example, phenomenal properties and epistemic normativity can play an irreducible explanatory role.

The proposed view is supposed to tell us how thoughts that draw on perception without descriptive or 'conceptual' mediation acquire the "kind of intentionality characteristic of thought" (Dickie 2011). In doing so, it tells us what perception lacks. Perceptual states have either "another kind of intentionality" (2011, ft. 15) or merely informational content in something like Evans' (1982) sense (Dickie 2015, pp. 148-150). Perceptually grounded thought represents an advance on perception in that it marshals these informational resources to arrive at singular content not attributable to perceptual states themselves.

The basic picture is that perceptual object files supply an unmediated perceptual link with objects in a way required by *non-descriptivism,* while *non-descriptivism* and *epistemic homing* explain why *content uptake is* true*.* Objects indexed to object files are 'referentially basic', in that capacities involved in referring to them in thought and language do not essentially involve conceptually mediated description. While judgments formed by uptake from a perceptual object file will typically inherit the files' property attributions, the anti-descriptivist condition on how the files work is also inherited by the belief. What matters is that the file is a reliable source of information about o, not that it accurate with respect to its properties across all possible circumstances. Meanwhile*, non-naturalism* permits a defence of *epistemic homing* that draws on the personal-level practical normativity that governs the cognitive skills involved in making oneself knowledgeable about things in one's environment. This, on Dickie's account, is tied up with a proprietary phenomenally conscious cognitive state, characterised as the mind's 'basic need to represent things outside of itself' (2015, ch. 3). So, while there's a distinctive phenomenological profile to perceptual acquaintance, it is not determined by the properties of objects of acquaintance.

Excluding *non-naturalism*, which she describes as a "framing assumption" (p. 13), Dickie defends all of these claims in detail. Rather than trying to do justice to the particulars of the account, my summary is meant to set up a discussion based on the following question: What is the relationship between the contents of perceptual object files and the contents of corresponding perceptual experiences? In doing so, I fit elements of Dickie's view together with the example

taken from Merleau-Ponty above and the previous sections' discussion and interpretation of the empirical literature on object files.

c) Perceptual acquaintance as epistemic skill

The core of Dickie's argument is a claim about what perception does for subjects who are motivated by a basic need to represent things outside of themselves: it offers a means to satisfy the need. While the claim that cognising subjects constitutively occupy this kind of motivational state is controversial, the characterisation of perception as a general way of satisfying its putative goal isn't. Theorists of perception since Aristotle have emphasised the importance of the fact that perception furnishes the subject with information about her environment, providing the subject matter for a significant part of her 'post-perceptual' cognition. On the face of it, these claims have an attractive symmetry: singular thought emerges where the need to represent and the means to do this meet in perceptual experience (cf. Dickie 2015, pp. 108-113). However we assess these general claims, we already have reason to see perceptual acquaintance as bound up with a proprietary kind of motivated activity. It's a consequence of the previous chapter's argument linking the unity of experience and agency that (unified) subjects have more specific *needs to represent,* determined by the intentions underlying the perceptual action space. Generally: intending to act on an object requires an agent to represent that object. So, we have independent reasons to explore an account of acquaintance as a means to an end or some ends – and this suggests an aspect of skill involved in meeting that (those) end(s). In this section I will focus on the connection between the motivational component of this formular and the putative skill involved in acquaintance. Differences that follow from anchoring this account either to a generalised cognitive need or to successive agential needs will come out in in the next section's criticism of Dickie's model.

To make this idea more precise, let's introduce a general notion of an epistemic skill:

> *Epistemic Skill*: A subject, S, is epistemically skilful with respect to a domain, D, iff a) S can exercise an ability, ϕ, governed by norms that track the truth and falsity of beliefs about D; b) ϕ is a way of forming beliefs; c) enough of the beliefs S does form about D are arrived at by ϕ-ing.

Epistemic skill departs in some ways from Dickie's official formulation of the equivalent claim,[94] but I think it capture what's essential to the motivation behind *epistemic homing*.[95] I've followed Dickie in formulating *epistemic skill* as a biconditional. (If you think there are kinds of epistemic skill that don't fit this schema, take it instead as characterising a core set of examples). I'll assume that at least for ordinary perceptually grounded beliefs, epistemic skill is a matter of reliable belief formation that meets some kind of 'anti-luck' criterion. I also adopt the notion of justification as non-lucky reliability from Dickie with no additional argument.[96]

A need to accurately represent things in a domain is what an epistemic skill with respect to that domain satisfies. An initial objection to this framing might note that 'need' carries specific connotations that are not always appropriate for a motivational state corresponding to a skill. Someone can be a skilled dancer, without being driven by a need to dance or in dancing to be successful only insofar as this need is satisfied.[97] What makes the need-skill pairing Dickie has in mind plausible is, on one hand, the domain in which the skill is exercised and, on the other, the kind of normativity that attaches to successful performance. In the first instance the domain is highly general, including all possible objects of cognition. The thought is that it's constitutive of being a cognitive agent that one is motivated to think about *something.* For the most salient cases of belief and judgement, this quickly gives rise to constitutive epistemic norms. For example, if you are not trying to be get anything right about anything you are assertoricly thinking about, then you violate norms that are internal to the activity of forming beliefs. Local deviations from these are possible, but if you just never tried to get things right, you wouldn't really be the sort of cognitive agent that believed anything.[98] For just such cognitive agents, it make sense to think of the motivational state behind this kind of norm-guided activity to be pervasive and urgent: just as creatures whose biology is organised around respiration need food,

---

[94] See Dickie (2015, pp. 90, 125-126; 2016). Here, as elsewhere, my restatement gives a précis of a more detailed claim, introduced over several steps. Some aspects of the whole picture that I think are not directly relevant to the current argument have been lost in compression.

[95] What's enough?

[96] This is a characterisation of justification, not a proposed analysis. However we should understand justification, I take it that it will at least entail that a justified cognitive attitude will not be arrived at *by luck*.

[97] The converse of this may also be true.

[98] This thought is closely related to the interpretationist framework defended in chapter (§2). But for Dickie the explanatory order between norms and psychological states is reversed: it is an essential feature of psychological states like belief and the need to represent that they are subject to norms of rationality, and this is why psychological agents can be interpreted in light of these norms. On my preferred interpretationist view, norms 'come first' in that it's because they capture real patterns in the behavior and psychological lives of subjects that it makes sense to attribute content bearing states to them.

creatures capable of thought need 'food for thought' (cf. Dickie 2015, p. 130). In the case of belief, this need is best met by forming true beliefs and avoiding falsehoods.

Whatever the status of these programmatic claims, if D is a subject's immediate physical environment, *epistemic skill* can be used to generate a plausible-looking explanation for her having epistemic skills in relation to it:

> *Epistemic object-perception:* S is epistemically skilful with respect S's environment partly because S's perceptual object files are governed by norms that track the truth and falsity of beliefs about the environment and enough of S's perceptual beliefs are generated by uptake from perceptual object files.

Here, much of the work motivating the general 'need to represent' can be done by appealing to the ecological considerations in the background of the empirical object file literature. If, as I've recommended, we understand the common target of personal-level and subpersonal psychological explanation to be the psychological capacities of subjects and agents (see §2.4), *epistemic object-perception* should follow easily from this subpersonal story and a naturalised version of *epistemic skill*. The claim is just that a subject's subpersonal capacities for visual object perception constitute a distinctive way for her to become knowledgeable about the perceptible properties of the objects in her environment.

Remember that, for Dickie, singular reference is a feature of the thought and talk of whole persons. Only they are directly subject to the kind of epistemic norms that guide the information marshalling strategies invoked in this story. Subpersonal processes are subject to the same epistemic norms in a derived way, by being tied into these processes. We've now seen that, at least paradigmatically, these are strategies for ensuring that beliefs about things in the environment are true. Since beliefs are cognitive states of persons, norms of belief formation are personal-level norms. What you are able to refer to on the basis of perception alone is fixed by what you are justified in believing on the basis of your perceptual input. And this is determined by the contents of your perceptual object files. So, there's something mistaken about attributing singular content to subpersonal representations in the early visual system considered in themselves. Reference to particulars is fixed by a whole agent's activity in accordance with norms that likewise only apply (or only directly apply) to whole psychological subjects. As a result, the story Dickie offers about how information gets worked up into singular representational content necessarily appeals to both explanatory levels – directly at the personal level and indirectly at the subpersonal level (Dickie 2015, pp. 96-99). It is irreducibly a two-level account.

Let's see how this story deals with the cases we've introduced already. First, we'll discuss Austin's decoy duck, then we'll return to *Maurice and Bertie*.

> *Decoy duck:* S sees a decoy duck. S's perceptual situation is such that cues from the duck fall within an overlapping range with cues that she would have received if the situation were the same, but there was a duck rather than a decoy in view.

The subject in *decoy duck* is aware of a perceptual object that belongs to an equivalence class that includes both ducks and decoy ducks. If she takes it she is seeing a duck, she will be wrong – but she might be nonetheless justified in this belief. After all, if the decoy *were* a duck, then we wouldn't usually hesitate to describe her as knowledgeable about it. Seeing a duck-like object in ordinary conditions is a way to be justified in a range of beliefs about that object. What falls within this range may be affected by context. The range might be narrower if, for example, S were on a visit to 'decoy duck county'[99]. But some of S's beliefs about o will always be justified by perception. If she is seeing a perceptual object, then S's means of justifying <duck> and <decoy duck> beliefs converge on the same environmental particular.

The possibility of justified beliefs with singular content makes possible the content of singular beliefs that do not share their hight grade of justification (e.g., <that is a duck> thought of a decoy in decoy duck county, or even <that is a toy boat> thought by a poorly situated perceiver). However, these beliefs will be 'weakly justified', since they are formed via an information marshalling strategy that is a reliable generator of non-lucky true beliefs about items in its domain (i.e., perceptual objects). This kind of weak justification is sufficient for singular content since it is incompatible with the non-existence of *some* perceptual object – that is, there must be something that the perceptual system is 'treating as an object' for the object file to be fulfilling its ordinary function. If the object file generates representations in the absence of any environmental input, then it will not be functioning as it should and so won't make its distinctive contribution to the epistemic status of downstream beliefs. This is a good result for the account. But so far, we have assumed that object files are at least systematically reliable. What if we try to do without this assumption? Depending on how the perceptual situation is further specified and how it is fitted into a wider epistemically relevant context, there may be variants of *decoy duck* in which everything the subject believes on the basis of her perceptual contact with an object would fail to be even weakly justified. We might worry that there will be cases in this

[99] See Goldman (1978). For recent discussion in connection with disjunctivist theories of perception, see for example, Neta (2007), Schellenburg (2017).

range in which the subject fails to be epistemically skilful, but it would be counterintuitive to deny singular content to her perceptual beliefs.

The constraints on reference imposed by *epistemic skill* are less stringent than someone expressing this worry supposes. We can start to see why by thinking about what all members of the equivalence class of duck-like perceptual objects that could be placed into the subject's perceptual situation have in common. Trivially, any such object would make a difference to the causal structure of a part of the environment that she has in view, combining with other elements of the perceptual situation to give rise to a characteristic range of appearances. It's only by radically changing the situation that the subject is in that we can construct a scenario in which she is able to form beliefs with intuitively singular content that are unjustified with respect to any of their objects properties or features. These changes bring the situation outside of the competencies of the epistemic skills deployed in visual perception. Here we do have to say that the subject is not perceptually acquainted with the putative objects of her perception and her perceptually grounded beliefs. But it's unclear that in departing from the bounds set by competencies exercised in ordinary perception, we still have reason to ascribe singular content to perception or belief. Dickie's suggestion here is to apply the standard contextualist response to the epistemic worries generated by this kind of case (the subject is not rationally required to eliminate the possibility that she is perceiving a virtual environment) to the semantic worry too. So, just as someone in a virtual environment might be unluckily false in (almost) all of her beliefs, when someone tries to bring their epistemic capacities to bear on appearances generated by a virtual environment the singular content of their 'perceptual' beliefs will be empty. Moreover, Dickie has a simple diagnosis for any lingering, intuition-based doubts. In this case, a subject's beliefs might have descriptively mediated singular content (e.g., ' the object causing my experience is P'), which depend on cognitive capacities that are not required and typically not deployed in genuine cases of perceptual acquaintance.[100]

Now let's consider how this account can help us understand the more complicated scenario described by Merleau-Ponty and the adjusted version of it presented above.

> *Maurice:* Maurice sees the ship in front of the trees. The mast is fused to the ship, where it should be. He distinguishes the boat as a whole from its background and is visually aware of the stand of trees, but not with the same degree of attentional focus. For example, he is not aware of each tree individually.

[100] Another response, reject 'actualist' account of singular content. Hansen & Rey (2016)

This is the good case. A subpersonal explanation of Maurice's situation, along the lines suggested above, attributes to his visual system an object file which successfully picks out the whole boat and assigns properties to it in a way that matches the way the boat is and how It looks (viewed from *here,* in *these* lighting conditions, etc.). The boat is thus at the focus of an information gathering process, in which converging streams of sensory information processed across distinct 'pre-attentive' processing areas are assigned to their common environmental source. The packaging of information on the boat's objective properties and some features of the perceptual situation itself (transient surface reflectance, perspectival size, etc.) in the object file allows for direct, attention-mediated pick up into personal-level belief formation. Moreover, it looks like we can understand the phenomenal 'instability' of the earlier scene on the model of the epistemic normativity guiding Maurice's information gathering procedure. There is a kind of 'normative magnetism' (cf. Williams 2020) to the correct scene: perceptually referring and getting things right about a scene seem to come together. This reflects the fact that the perceptual channel after the shift is a reliable source of more (and more precise) content about the ship with the mast than without it.

Meanwhile, the low reliability of judgements that Maurice could make about individual trees behind the ship corresponds to an inability to form basic perceptual demonstrative thoughts about them as individuals (cf. Dickie 2015). Conceptually unassisted thought about a particular tree in the unattended group would get things right about it rarely and luckily. The subpersonal explanation for this is that there will be no corresponding object file for cognition to draw on – the visual system does not work in a way that allows for Maurice to index and track everything in the scene at the same time.

What about the bad cases? Recall that Bertie$_a$ succeeds in getting Maurice's linguistic referent right but *perceptually* fails to pick out the boat in a way that includes the mast. Meanwhile, for Bertie$_b$ a more basic perceptual failure ramifies into linguistic interpretive failure.

> *Bertie$_a$:* Bertie sees the ship in front of the trees. The mast is detached from the trees. He distinguishes the hull of the boat from the background and is visually aware of the stand of trees, but not with the same attentional focus. He is not aware of each tree individually, nor is he aware of the mast as something distinct from the stand of trees.
>
> *Bertie$_b$:* As with *Bertie$_a$*, but now Bertie picks out at least one individual from the stand of trees, but this is – in fact – the mast of the ship. So, he is aware (illusively) of the mast as something distinct from the boat and (veridically) as distinct from the other trees around it.

Epistemic object-perception supplies principled backing for the intuitive verdict that Bertie$_a$ picks up on the correct referent for Maurice's 'that' while Bertie$_b$ fails. Both pick out an object that is in fact continuous with the boat referred to in *M* – either the hull or the detached mast. But only Bertie$_a$ is able to arrive at a range of well-grounded judgements about the boat by picking up attributes assigned to it by the relevant file, although he's not yet able to share Maurice's <o is tall> judgment. A possible subpersonal explanation for this is that in Bertie$_b$'s case the detached mast has been indexed to its own object file, maintained against the pull of epistemic gravity towards the complete scene. If this *is* what's happening, it will in principle be open to mechanism-targeting investigation. This would reveal that although the *environmental* causes of their perceptually grounded judgments as the same across cases (down to a near-identical 'retinal image' projected by the light array from the scene), only Bertie$_a$'s perceptual link with the object indexed to a perceptual object file counts as a deployment of his epistemic skill with respect to the same thing as Maurice.

Having locked onto the right target, it's natural to imagine Bertie$_a$ undergoing the change already experienced by Maurice. This might be via 'bottom-up' processes that segregate the boat's parts from the perceptual background, or via 'top down' effects of the conceptually mediated judgement that Bertie$_a$ is prompted to entertain. Keeping Bertie$_b$ in the frame, we can easily imagine him correcting *both* mistakes at once ('Oh, the boat? Now I see it…'). But it's very hard to imagine him noticing that what he took to be a tree was in fact the mast of the boat and nonetheless persisting in the belief that Maurice was pointing to *some tree or other* in the rough area he was pointing at, behind the now noticeably taller boat.

If this is the case, we should want to explain a capacity to refer to things in the environment partly in terms of a capacity to get things right about them: to fix onto, track and correctly attribute properties to objects. Dickie's account offers a way of integrating these intuitions with a plausible interpretation of the object files literature: files index objects and thereby enable perceivers to know the intrinsic properties of those objects. This knowledge *can* be used to formulate an individuating description, but it relies on the effective use of more basic capacities for indexing and tracking.

2) Object Files and the Subpersonal Conditions of Acquaintance

a) perceptual attention and isomorphic explanation

I've said that naturalistic versions of this proposal are unavoidably committed to explanatory isomorphism. It's time to support this claim. To do this we need to bring perceptual attention, which has so far been in the background of the picture, into focus. This is because the skills

deployed in *epistemic object perception* are skills that the perceiving subject exercises in her deployment of attention. Subpersonal mechanisms for perceptual object files supply the material out of which singular content is constructed, but attention is what brings them together. It's a crucial feature of the view (on Dickie's presentation and on any naturalised alternative version) that we should understand attention as what Watzl (2011, p. 849) has called "the mental activity of structuring the stream of consciousness." For the account to work as Dickie means it too, subpersonal explanations of this activity must be isomorphic. And, as we've seen, isomorphic explanation of a phenomenon in terms of internal mechanisms precludes naturalistic routes to externalism about its phenomenal or psychological properties. If naturalism is assumed, the isomorphism debate overlaps with the internalism/externalism debate.

First, let's distinguish between personal and subpersonal notions of 'attention'. At the personal level, perceptual attention is a pervasive feature of people's psychological lives and of how they interpret those of others. Capacities for 'attention reading' are some of the earliest developing and most basic forms of social cognition, and there is evidence that they are the most widely distributed among non-human animals (Call 2009). Since attention can be both intentionally directed towards and spontaneously captured by objects, understanding a subject's attentional state often involves taking account of both endogenous (subject-to-object) and exogenous (object-to-subject) causal dynamics. Appreciating these is at the basis of attributing more sophisticated cognitive attitudes and psychological activities. Notably, these include the perspectival interactions at the focus of the current argument (e.g., noticing, observing, listening). Many of the distinctions captured in ordinary-language terms for these turn on the relative weighting of endogenous and exogenous dynamics, and personal-level explanations of what an agent intentionally does implicitly or explicitly evoke corresponding attentional states. Attention, then, situates subjects and agents at the centre of a network of causal interactions modifying more basic sensory and motor capacities and facilitating sophisticated cognitive achievements based on these (Hurley 2008). Appeals to attention in personal-level explanation are attempts to capture aspects of this complicated structure, played out between agent and environment.

But 'attention' can also play a subpersonal explanatory role – as, for example, in Triesman's early models of feature binding. Here 'selective attention' is the name of the mechanisms that binds features together and maps feature clusters to locations in retinotopic and environmental space (Treisman 1969; Humphreys 2016). On other subpersonal theories, attention is associated with selection of salient information processed locally in neurocomputational subsystems for 'global

broadcasting' (Prinz 2000), or with a filter on pre-attentive perceptual processing through which a small number of available proto-objects pass to become targets for object-file systems (Shinn-Broadbent 1959; Cunningham 2008). More recent 'top-down' models posit processes of 'biased competition' (Desimone 1998), associating attention with patterns of feedback and feedforward between high-level representations underlying cognitive states and 'sensory' representations in early perceptual areas (for recent reviews of the cognitive scientific literature, see Mole 2009; Keya et al 2017; for philosophical discussion, see Wu 2012; Watzl 2011, 2023).

Some naturalistic isomorphic programmes identify the personal-level phenomena with a subset of mechanisms of this kind. Prinz (2000) defends a radical proposal along these lines, claiming that the contents of perceptual experience are all and only the contents of attended perception, and the contents of attention are contents selected for global broadcast by attentional mechanisms. Alternatively, isomorphic assumptions can be used to motivate eliminativism about the folk notion, or about associated personal-level constructs. In an influential early review of neurocomputational models of attention, Allport (1993) defends this response, noting that

> the heterogeneity and functional separability of different components of spatial and nonspatial attentional control prompts the conclusion that, *qua causal mechanism, there can be no such thing as attention*. There is no one uniform computational function, or mental operation (in general no one causal mechanism) to which all so-called attentional phenomena can be attributed. (Allport 1993, p. 203; emphasis added)

If explanatory isomorphism is assumed, then failure to assign the 'mental operations' of attention to a causal mechanism entails that the personal level phenomenon is not a unified cognitive kind. Allport suggests that investigating the 'mechanisms of attention' is like investigating the 'mechanisms of thinking' – the folk notion used in personal level explanation and interpretation is too heterogenous to support mature cognitive scientific enquiry. More recently – and succinctly – Hommel and colleagues sum up a similar argument. What accounts for the variety of explanatory roles, experimental paradigms and proposed mechanisms that contemporary cognitive science groups together under the label 'attention' is that (James' famous statement to the contrary notwithstanding) "no one knows what attention is" (2019, p. 2288).[101]

[101] James (1890) on attention as mental highlighting

Taken at face value, this assessment is overly pessimistic. The ways in which psychologists have operationalised attention are targeted at a much more unified pattern of psychological and behavioural dynamics than could plausibly be attached to an unrestricted notion of 'thought'. If we assume the conclusions of chapter (§1), this should suggest to us that a non-isomorphic strategy is appropriate (cf. Watzl 2011). Such an approach proceeds by specifying a set of capacities and functions for attention manifested in the behaviour and psychology of the whole agent and then investigating the constitutive contribution of a range of subpersonal elements, including mechanisms like those identified above. So, for example, conscious attention might be associated with a distinctive phenomenological and behavioural profile at the personal level, which are in part explained by the neurocomputational role played by multiple subpersonal mechanisms. Neither phenomenological properties nor the role that attention plays in the psychological life of the subject need be assigned to attentional mechanisms, which coordinate with other subpersonal elements to jointly constitute these personal level phenomena. In general, if a proposed model of perceptual attention recognises the possibility of non-isomorphic interlevel explanatory connections, it can distinguish between personal-level conscious attention and 'attentional mechanisms' appealed to at lower levels. With this distinction in place, we might then put it to explanatory work in an overall two-level picture.

Alternatively, granted an *ontological* distinction between personal and subpersonal levels, we can ascribe an autonomous causal-functional role to personal-level attention. In something like the way a statue has modal (as well as aesthetic, economic and social) properties distinct from the particular parcel of matter that materially constitutes it, personal-level attention while *depending on* underlying mechanisms might impart psychological or semantic form onto their processing that goes beyond their informational properties considered in themselves (cf. French & Philips 2023). For example, consider the role of personal-level attention in Campbell's model of perceptual acquaintance:

> Conscious attention to the object has in effect to provide an address for the thing, so that the right information-processing procedures can operate for verification or action […] [It] identifies the object by the binding principle for that thing; the principle that the visual system uses to put together all the information relating to that object as all concerning one and the same object. (Campbell 2002, pp. 253-254)

Campbell's 'conscious attention' is a personal-level capacity that stands in a guiding or directing relation to the subpersonal mechanisms that create and maintain perceptual object files. As a

distinctive personal-level state that stands in for the subject in orienting the subpersonal resources that make her experience possible, it plays a similar role to Velleman's (1992) 'full-blooded' intentions, fusing event-causal and normative roles (§3.3c). On this picture, attending to an object is a way of causally intervening in subpersonal mechanisms. Elsewhere, Campbell describes attention as "what defines the target of processing for the visuomotor system, and thereby ensures that the object you intend to act on is the very same as the object with which the visuomotor system becomes engaged" (p. 55).

On Dickie's version of this story, perceptual selective attention realises the subject's epistemic skill in responding to the outputs of feature-to-property computations in the early visual system. Although these are internally realised by early (i.e., 'preattentive') perceptual mechanisms, the proposal is not that the content of experience simply reflects the representational states of a specifiable subset of these mechanisms, along the lines suggested by Prinz (2000). Instead, the explanatory mapping in her model is between elements in the domain in which the subject exercises her attentional ability and elements in the domain on which subpersonal mechanisms for attention operate:

> 'Perceptual selective attention' is the highlighting of elements of perceptual input for additional processing […] At each moment, the part of your visual field processed to this higher resolution is the part in the spotlight of your visual selective attention. [The] property-from-feature calculations that give rise to property-level appearances (which in turn give rise to perceptual demonstrative beliefs) are concentrated within the attentional spotlight. (p. 118)

If we think of the visual field as something distinct from 'perceptual input', it's unclear here which of two pictures the metaphor of the attentional spotlight is supposed to evoke. Does attention cast its beam over the environmental objects of acquaintance or over the internally realised appearances that make up the visual field? We require a tight correspondence between the two to say that attention 'passes through' the appearances to its objects – then attention to the outputs of early perceptual processes (under the right conditions, granted the cooperation of the environment) just is attention to the distal environmental objects.

This reading is further supported when we look at what Dickie says about the elements that fall outside of the attentional spotlight. These "do not produce fine-grained perceptual experiences" (p. 119) that support reliable direct judgments. Consequently, she continues,

> you will not report that something in unattended peripheral vision is [for example] beige and cubical because property-from-feature calculations outside the attentional spotlight do not deliver this much fineness of grain […] they are not reliable. (ibid.)

This is plausible if we think of the perceptual field as the reportable content of perceptual representation fixed at a moment. It is less plausible if we think of the objects of perception as those that are perceptually available to a subject in an ordinary episode. Think back to *Maurice and Bertie*. Although individual trees that form the background to the boat are not available for basic demonstrative thought, it is counterintuitive to claim that a subject will be unable to report their features unless 'aided' by conceptually mediated description. For this to be the case, we have to hold the scene and the attentional spotlight fixed – replicating the conditions of some laboratory investigation of perceptual attention but diverging from the conditions ordinarily targeted by personal-level explanation and interpretation. In ordinary conditions, subjects are consistently reliable about more than they attentively focus on *in any given moment* because they can redirect attention within a shorter timeframe than it takes to formulate reports. (Sergent 2018). It requires non-standard conditions and manipulations, like those that elicit change-blindness effects to produce the unreliability that the quoted passage alludes to. If acquaintance precludes the kind of pragmatic access that ordinary subjects have to things around them in the course of daily activity, this can only be because of a background interlevel isomorphism assumption.

Thus, on the model offered by Dickie there is an explanatory isomorphic mapping between a) *appearances*, available to the subject and on the basis of which she is able to formulate singular thought and judgment about external objects, and b) *outputs from the early perceptual system*. It should be noted that 'appearances' are introduced into Dickie's discussion in traditional first-personal terms, as features of a scene presented to perceptual consciousness. After they're linked to corresponding subpersonal computational events and processes, however, a further set of correspondences is required. These link (a) and (b) to a third domain, constituted by the environmental objects of perception and their behaviour:

> This [acquaintance supporting] attentional link involves conceptually unaided property-from-feature calculations – calculations that require a characteristic coherence in feature-level appearances and a corresponding kind of coherence in the behaviour of the attended thing. (p. 128)

‘Feature-level appearances’ are aspects of the phenomenal character of perceptual experiences. They contrast with ‘property-level appearances’, with the former corresponding to the aspects of objects’ appearances that change across perceptual constancies. Property-level appearances, meanwhile, are constant. For example, a visual experience of a spinning circular disk will exhibit property-level appearances for ‘roundness’ and ‘turning’ and feature-level appearances for more or less flat ellipses (pp. 116-119). The correspondence between object properties and property-level appearances is an upshot of feature-to-property computations that track cue-from-object dependencies back to the object. Where the perceptual situation is unusual, this can give rise to illusory appearances. But in good cases, attended environmental object corresponds to a distinctive pattern of appearances, which allow for property-level appearances appropriate for justification-securing uptake into perceptual belief.

So, the overall picture is as follows. Good epistemic conditions are those that foster reliable, non-lucky perceptual judgment. They eliminate the epistemically relevant conditions in which objects and appearances fail to match, so that when these conditions obtain there is a three-domain correspondence between environmental objects and their properties, phenomenally conscious states of the subject and informational states of facilitating mechanisms. States of acquaintance supervene on these states, which are the constitutive goal of the subject’s attentive perception if her environment. But, although the content of the states depend on the appearance-environment relations targeted by isomorphic explanatory mappings, the way good cases are specified gives no additional psychological role for objects to play than occurring in the immediate causal histories of appearances as these play out, moment to moment.

b) The Accomplishment Model of Acquaintance

We’ve seen that Dickie’s model of perceptual acquaintance combines traditional and novel elements. What’s new is its focus on perceiving subjects’ skill and agency, which casts the acquaintance relation as something the subject plays an active role in forging. Cognitive agents are constitutively “engaged in the activity of marshalling information into bodies of belief about particular things” (Dickie 2015, p. 307) – being perceptually acquainted with something necessarily involves corresponding activities of getting and staying acquainted with that thing. The more traditional elements of her model involve the way this goal is specified and the domain in which these activities are is manifested. For Dickie, getting acquainted involves a disembodied cognitive skill exercised by the agent in directing the attentional spotlight over a perceptual field composed of features and properties (see, e.g., 2015, p. 286). The goal of this activity is categorical knowledge of the objects of acquaintance. So, on one hand, acquaintance is a kind

of goal state – perceptual knowledge or belief – which these activities target and maintain. On the other, it is the marshalling activity itself. The state of being acquainted with an object involves something like what Ryle called the "more or less protracted proceedings" involved in "keeping a secret" or "holding the enemy at bay" (Ryle 1949, p. 149 ft.). This claim is hard to motivate when the model for acquaintance is introspection of decontextualised states of perceptual consciousness. But it finds a natural home in the context of the previous chapter's claims about the durational character of perspectival experience. Nonetheless, I'll say a little more to bring out this connection below.

Notably, it's the traditional rather that the novel aspects of this picture that are in tension with phenomenal acquaintance. The most natural solution for externalists here is to *change the domain*. More precisely, the solution would involve denying that there are the envisioned correspondences between property- and feature-level appearances and neurocomputational components of feature-to-property calculations in the visual system (and, more generally, in other distal perceptual areas), and so denying that the isomorphic explanatory framework assumed by Dickie is appropriate. On this proposal, appearances map to environmental object, but their *transparency*$_{Ext}$ is not explained by positing a system of one-one correspondences that emerges under epistemologically good conditions. Rather, good conditions are just those in which subject-level target activities that supervene on these environment involving states are met, even if there is no synchronic one-one mapping from object property to computational state (to occurrent appearance). After all, in good cases we already have the personal-level description (e.g., 'S notices', 'S recognises', etc.) – and if it's apt, *we have the truths it expresses* (cf. Dennett 1969, p. 19) – before we introduce a system of correspondences like this. When a subject perceptually attends to some object, we're already justified in concluding that *some* combination of subpersonal elements cooperates to open and maintain an attentional link between person and thing. And we've already rejected the claim that the *only* way to understand what is happening at this lower level of description is to reinscribe on it the same set of relations as we used to characterise the personal-level phenomenon.

In making good on this suggestion, it will be helpful to have a name for the novel aspect of Dickie's view, which it shares with some other recent proposals (see Soteriou 2013; Ward 2016a, 2023; Raleigh 2021). Call these *accomplishment models* of acquaintance. My use of the term 'accomplishment' is inspired by Vendler's (1957) taxonomy of the 'time schematas' presupposed by different categories of verb. For psychological phenomena, these correspond to what Steward (1997) has called the 'temporal shape' of psychological 'continuants' (roughly, states

and processes) and 'occurrents' (i.e., events).[102,103] One broad category of these count as good answers to the question 'what are you doing?' On a first approximation, *running*, *listening* or *calculating* belong in this category, *knowing* or *loving* don't. Call members of this category 'activity verbs'. Vendler notes that many activity verbs (and phrases constructed from them), are directed towards an end or goal – these are verbs for 'telic' activities (cf. Hornsby 2013). 'Running a mile' is an example, as is 'calculating the required quantity of water' or 'drawing a circle'. When a telic activity is successful – the circle is drawn, the quantity is calculated – these verbs attribute *accomplishments* to their subjects.

It's a familiar point that accomplishments and other activities occupy time differently from states (Steward 1997; Soteriou 2013). Someone might learn that the fastest land-animal is the ostrich, that the knack to a good somersault is in the pressure applied on launch, or that one should generally develop knights towards the centre of the board, and they might continue to know these things for their whole life afterwards. But this does not mean that there is or must be some corresponding activity of 'knowing' that they are engaged in from that moment onwards. On the other hand, it also seems that some things that a person does don't occupy any significant interval of time at all. Winning a race, delivering checkmate or deriving the result of a calculation might be achieved 'in an instant'. Accomplishment models present acquaintance as the constitutive target of a corresponding activity of *getting acquainted* that unfolds over time and is structured towards this goal. Consequently, the activities that accomplishment theorists appeal to in their explanations of acquaintance are non-static and non-instantaneous.

The following biconditional is true of members of the broad class of 'activities':

> *Active Continuity:* For an activity, ϕ, and a subject, S, S ϕ's at a time, t, iff, there is an interval, t_1 - t_n, containing t, such that S is ϕ-ing from t_1 to t_n

[102] Steward's (1997) work has been seminal for a recent tradition (Soteriou 2013; Crowther 2009; Hornsby 2013) reviving interest set of questions in the ontology of mind about relations between different kinds of 'mental continuants'. These were central for Ryle (1949) among others (see, e.g. Dennett 1969 – Anscombe chapter), but received less attention with the increasing centrality of 'states' (cognitive and representational) to mainstream discussion. Soteriou (2013) brings these considerations to bear in a series of arguments against representation-based views.

[103] In subpersonal explanation, 'activities' can naturally be replaced by 'processes', so we might adopt this term generally (cf. Crowther 2009). When discussing personal-level agency, however, the term 'activity' is preferable since this allows us to more easily draw connections (and mark distinctions) between sensorimotor interactions and 'actions' as discussed in the previous chapter.

There's an intuitive sense in which a non-telic activity like running 'takes up' or 'occupies' an interval of time in a uniform way. We can capture this by supplementing *active continuity* with a further condition:

> *Homogenous Continuity:* If S is ϕ-ing from t_1 to t_n then for every t from t_1 to t_n S is ϕ-ing at t

Something of which *active continuity* is true but *homogenous continuity* in not is, according to Vendler, an accomplishment. This specific difference marks out accomplishments in virtue of their goal-oriented, telic, temporal structure. If, halfway through running a mile, someone gives up, then although she has been *running* for some interval of time, she hasn't been *running a mile* at all. Accomplishments are necessarily non-homogeneous.

Neither *active continuity* nor *non-homogeneity* are obvious properties of the acquaintance relation. Recall that Russell characterised acquaintance as a mode of *presentation* (Russell 1911, p. 108). Being presented with an object is not an accomplishment. Russelian acquaintance, like love at first sight, is something that *befalls* you. Based on introspection, this may seem faithful to many – but certainly not all – perceptual experiences. Hearing the waves wash up onto the beach or watching the sun slowly sink beneath the horizon don't seem to be things that one takes an active role in doing. Moreover, cases which do plausibly have an effortful, agential component (e.g., bringing the beached ship more clearly into view) seem to be non-basic cases, apt to be reduced to complexes of these simpler experiences. The thought is that perception, in the basic case is passive.

Nonetheless, two considerations speak in favour of the accomplishment theorist's approach:

a) *Unfolding experience:* perceptual experiences of objects track changes in their perceived features and properties over time in a continuous way. Hearing a melody, seeing a duck glide across a pond is something that fills the interval it occupies: it exhibits active continuity.
b) *Causal processes:* the mechanisms of perception constitutively target representational states, receiving input from the world outputting representational vehicles that are interpretable via appeal to norms of environmental appropriateness. If, as previous chapters have argued, content attribution depends on how these input-output mechanisms are fitted in to a wider context (contents ascribed to outputted vehicles are not intrinsic to those vehicles) then

> interpretation and content attribution needs to consider at least the recent local history of the processes: they are non-homogenous.

Of these (a) focuses on personal level and (b) focus on subpersonal aspects of perspectival experience. While Dickie gestures towards (a) (2015, p.136 ft. 35.) her focus is more on considerations mentioned in (b). This is because, with her focus on a static goal of acquaintance (perceptual belief) her interest is in how this goal state is produced. But we needn't, and I think shouldn't, think of experiences as horizontally constituted by successions of states. And if we don't do this, we can use the accomplishment model to explain phenomenal acquaintance realised by perspectival experience. As we've seen, accomplishments can be constituted by other continuants – paradigmatically processes and activities, but also states (for example, winning a game of chess involves achieving a series of board states). The heterogenous character of accomplishments frequently resists the attempt to align their temporal parts with their constituent elements. It's only in the case of simple accomplishments that we can assign one-one correspondence between, for example, the temporal parts of an athlete's *winning the race* and of her *running* in it. So, what the accomplishment theorist needs to establish is that acquaintance does not have this simple structure.

To adopt the strategy I'm proposing, phenomenal acquaintance theorists need to provide an account of how interactions between neurocomputational mechanisms for feature binding, tracking and object-property attribution link up with environmental particulars in a way that explains the phenomena highlighted in (a). On this kind of two-level interdependence view, the suite of cognitive abilities deployed by the subject in attending and responding to things in her environment neither reduce to nor epiphenomenally 'emerge' from any of the elements of the underlying subpersonal story. It is because the subject is perceptually situated as she is that the underlying mechanisms and interactions are coordinated, but equally her being so situated is not a further metaphysical fact. The activity of the mechanisms in their specific spatiotemporal context *just is* the subject enjoying perspectival experience. If arguments against the reduction of bodily agency to basic action are been successful (see §3.4), we already have a set of paradigm examples: complex bodily activities.

3) Getting Acquainted: Agential Skill and Perceptual Objects

a) Perceptual Acquaintance as a Sensorimotor Accomplishment

The sensorimotor models of perception are ideally placed to provide this account. Recall that sensorimotor accounts of perceptual experience make two characteristic claims. At the personal level, sensorimotor theorists argue that perception and action are constitutively

interdependent. This is the claim that chapter (§3) defended and put to work in an argument for perspectival unity. The result was the personal-level notion of an episode of perspectival experience. Environmental objects play a central role in perspectival episodes, since the explanatory norms we use to ascribe contents to the perceptions and intentions that constitute them rely on their subjects being situated in environments composed of objects. (I will assume these results in setting out the argument for the sensorimotor acquaintance view).

The second claim is about the subpersonal mechanisms of perceptual experience. Here's a general formulation (cf. Hurley 1998, ch. 10; O'Regan & Noë 2001):

> *Subpersonal interdependence:* The content of a perceptual representational state, P, is a function of a) sensory input b) motor output c) prior perceptual states d) neurocomputational processes that link (a) and (b) to transformations of (c)

The personal and subpersonal components of the sensorimotor picture reinforce each other. If a perceptual state is partly constituted by potential or actualised bodily action, then the constitutive role assigned to motor output in *subpersonal interdependence* follows trivially (Hurley 1998, ch. 10; Kirschoff & Kiverstein 2019). Remember, the claim needn't be that every perceptual state entails some motor output. In principle, the schema can accommodate 'null' perceptual outputs as long as there are systematic relations between inputs and outputs that *sometimes* require other values for the output to be specified (see §2.4). But it's also worth remembering that there being no appropriate personal-level action attribution does not amount to there being no motor output (Engel et al. 2015). The key point is that the informational content available to the perceptual system is sometimes as much a matter of output-input causal loops that run through the environment as input-output transformations internal to the system. Since the environment is composed of objects, and these objects are the targets of a very large proportion of the subject's intentional actions, this means that objects enter into subpersonal interdependence implicitly as mediating terms between output and input. By entering into these perception-constituting causal interactions, objects come to be part of the subpersonal basis of corresponding experiences.

Some accomplishment acquaintance theorists endorse this approach. For example, Raleigh, motivates a naïve-realist style 'relationalist' account of the phenomenology of shape perception by appealing to aspects of perceivers typical range of action with respect to objects' shapes:

> [T]he *way* or *manner* in which the object's shape and the perceiving subject's viewpoint *combine* in an experience, so as to give rise to *this* specific

> phenomenology, is non-mysterious and intelligible because both the object's shape and the visual system's position are familiar macroscopic features […] that it is (in general) easy for us to *vary independently and at will* whilst consciously experiencing the results. That is why we all have an implicit, practical and *sensory* grasp of the relationship—a relationship that is explicitly described by the theory of perspective—between the object's intrinsic shape, one's viewing position and the resulting shape appearances. (Raleigh 2021, pp. S4098-4099, emphasis in original)

The target here is the same kind of shape constancies that were at the focus of Dickie's account – for Raleigh, these are the ideal 'parade-case example' for the kind of agential non-isomorphic account he's discussing (see p. S4094), since the purported objects of acquaintance (shape property instances) have an obvious bearing on bodily action (cf. O'Regan & Noë 2001a,b). Sensorimotor accounts of these are promising, Raleigh suggests, because there is a range of everyday capacities for bodily action which can provide the implicit basis for 'epistemic' skills deployed on objects. In this case, then, accomplishing acquaintance needn't call on more recondite abilities to direct the attentional spotlight – and if it gets on without these (or without *requiring* these) it looks like we have a personal-level means of gaining transparency$_{Ext}$ to objects without the kind of isomorphisms introduced by Dickie's model.

More ambitiously, Ward (2016, 2023) has suggested that combining relational and sensorimotor accounts of perceptual phenomenology yields a general framework for answering the questions Raleigh poses for object shape properties. That is, a properly worked out version of this framework could account for the full range of perceptible properties in something like the way Raleigh thinks it might handle shape perception. Here's Ward's prospectus for this account:

> Relationalists can explain the phenomenal properties of perception in terms of the mind-independent properties of objects and situations while providing a sensorimotor account of what the perceptual relation consists in—of how embodied perceivers are put in touch with mind-independent properties via their grasp of their place in the web of contingences linking sensation and movement. (Ward 2023, p.259)

On this proposal, sensorimotor acquaintance takes on the same level of generality as the traditional acquaintance relation. It is not, as it is for Raleigh, one way of being related to environmental particulars (perspectival grasp of their macroscopic features) or a relation to a special class of perceptible particulars (intrinsic shape, rather than colour). The sensorimotor

acquaintance relation is supposed to stand in for and explain all visual perceptual relations, and possibly all distal perception (p. 258) Consequently, Ward's framework aims to combine a naïve realist answer to the scope question with a generalised naturalising explanation of the acquaintance relation. While the appearance properties that constitute an experience's phenomenal character are fixed by the intrinsic properties of the perceived object, your "openness to those properties is explained by appeal to your grasp of the sensorimotor contingencies involved in your current relation to the [object]" (p. 259). The proposal is that the particular object-dependent components of phenomenology are explained by acquaintance, while acquaintance is given a non-isomorphic interlevel explanation in terms of mechanisms that represent and respond to sensorimotor contingencies *plus* the intrinsic properties of objects of acquaintance. These are the basis of the "inarticulate, embodied knowledge" (ibid.) that plays an analogous role in this account to *epistemic skill* in Dickie's.

Setting these claims out against the framework established over this and previous chapters, we now have a way to support some of them. This is because we've got a way to characterise both the constitutive goal and the goal-oriented activity that combine in sensorimotor acquaintance.

> *Perspectival Unity*: Perception situates subjects in an environment that provides them a temporally and spatially unified action space. Objects carry and transmit causal influence and are individuated by the ways they fit into the casual structure of the environment. Situating the subject in this way is the constitutive goal of the subpersonal mechanisms of acquaintance.
>
> *Agential Integration*: A subject is situated in an environmental action space partly because her perceptual object files pick out objects in a way that respects their individuating causal relations. This contributes to her being interpretable according to norms of appropriateness to the environment in which she is so situated.

Agential integration says that feature binding, object tracking and property attribution functions are coordinated towards the goal of *perspectival unity*. *Perspectival unity* says that agents have a basic kind of knowledge of the environments they are situated in – they know how to manipulate its causal structure. On this way of developing the sensorimotor proposal, epistemic skill is recast as a special case of agential skill. Getting acquainted, on this view is something the subject does 'out in the open' (Hurley 1998) – it is an accomplishment of ordinary bodily perception and agency (and so, in part, the subpersonal mechanisms underlying these capacities).

This allows us to flesh out the two-level account of the interactions evoked in the above characterisations of *object-appearance* (§4.1). At the personal level, a perspectival experience is horizontally constituted by a pattern of perceptions and intentions converging on one or more objects. This pattern has a structural unity that is captured by *normative unity principles* (e.g. holism, environmental appropriateness) so that intentional and perceptual contents broadly cohere with each other and are intelligible against the situated psychological history of the agent. These patterns are targets for subpersonal explanation. Subpersonally, the structural unity of the experience is explained and justified it terms of the causal unity of the objects of experience. Environmental objects are mechanisms for transmitting causal influence, which explains their systematic contribution to cue-from-object dependencies that determine the patterns of input they provide to perceptual systems. Given the role these have in generating changes in the causal structure of the environment, these systems themselves help to shape the dependencies that determine the range of cues that are available to them. Object files integrate the results of feature-to-property calculation on patterns of input that their earlier processing has played a systematic role in shaping. Part of what underlies the reliability of the information stream available for the subject to 'pick up' from an object file-mediated perceptual channel is this kind of output-input interdependence at the subpersonal level.

Perspectival experiences are constituted by a stretch of perceptual activity on the part of the subject. It fills or unfolds across the interval it occupies, and this interval is determined by the subject-object interactions that the subject's activity partially consists in. These require the presence of an object, which in good cases is an appropriate target for an object file. Under ordinary conditions, appropriate objects sustain the activity of object files while inappropriate objects fail to. So, at the subpersonal level, these environmental object-subject interactions can be described as patterns of causal interaction that provide the right kind of input to object file mechanisms. By broadening the temporal window at the personal level, we make space to understand the subpersonal content that contributes to the subject's perspective on her environment. Since both external cue-from-object dependencies and internal input-output transformations are components of the overall causal structure on which the unified perspective depends, they constitute the real patterns in subject's psychological lives that are the target of personal level interpretation. The whole system composes a subpersonal 'dynamic singularity' (Hurley 1998), out of which elements can be abstracted and separately analysed as perceptual objects, representational mechanisms, motor- or sensory inputs, and so on. These abstractions allow us to get an explanatory grip on the behaviour of the system, which is exploited in subpersonal explanation. Nonetheless, personal-level interpretation continues to impose

constraints at lower levels. Content assignments to subpersonal mechanism components must make sense against the attributed goals and interests of the subject. Purported singular content of perceptual property attributions depend on the context of wider processing and the systems actual environmental situation. Generally, subpersonal explanation depends on personal-level interpretation, but personal-level phenomena ontologically depend on the coordination of elements that subpersonal explanation reveals.

Object properties contribute to the shape of these interactions. The singular contents of perceptual object files index the particular objects responsible for these systemic effects while the properties attributed to them capture aspects of those effects. Each perspectival experience is an unrepeatable episode in the psychological history of the agent, but episodes can be typed together by reference to shared structure in object-subject interactions. Subpersonal explanation enables greater articulation of these types and can reveal the basis for type-similarities that appear puzzling at the personal level. A subset of the regularities by which we type experiences is made up of those that make a difference to their phenomenal character. Phenomenal character is determined by the content and structure of an experience. Individuating an experience by its structure is locating it in relation to the other experiences that make up the experiential part of a subject's psychological history. Objects generate individuating structure of this kind by providing a focus towards which multiple cognitive and perceptual processes are oriented. They determine the temporal window in which content-determining subpersonal explanation of these processes can isolate mechanism components and attribute contents. This means that objects determine natural units of experience that can be the target of subpersonal explanation.

On the content side, subpersonal representations of object properties are both the causes and effects of the interactions that play out within these intervals. Attributing content allows researchers to mark the systematic contribution made by distinct subpersonal mechanisms and mechanism components to the subject's overall contribution to these interactions. In general terms, just as the properties of an object determine the overall contribution it makes to the causal structure of the environment, a perceptual or agential mechanism's representation of those properties makes sense of the agent's capacity to respond to and act on the object in a way that respects its causal relations – most saliently, these are its relations to the agent herself.

Some researchers have proposed explicit representations and dedicated representational systems that target these interactions. For example, Hommel's (2004; Camus et al. 2018) 'event file' construct is supposed to represent object-agent interactions in a format similar to standard

object files. Meanwhile, Cisek's (2007; Pezzulo & Cisek 2016) 'affordance competition hypothesis' aims to replace object files with constitutively action-oriented representational mechanisms that subject-object relations of this kind. But even within the mainstream object file framework, we can make space for a special role for object representations. Since so much of the causal structure of ordinary environments is captured by representing the objects in those environment, conditions of success for perspectival unity are closely tied up with the successful representation of objects. Object representation is a constitutive goal of the processes underlying perspectival experiences; when object-representing mechanisms are functioning as they should (in normal environments) the subject's experiences take on a characteristic phenomenal shape, that is captured and reported in perceptual singular content attribution at the personal level.

But while representational content allows us to capture the role of objects in phenomenal character, it's clear that we cannot read phenomenology off the content attributed at either personal or subpersonal levels (cf. Campbell 2008). Some contents are strongly associated with phenomenal differences, while others are not associated with strong or systematic effects. The sensorimotor acquaintance view offers the following explanation for this. Some properties and features are more relevant than others to predicting and explaining an object's causal relations to its environment on the timeframe of the object-subject interactions that constitute perspectival experience. Low-level features (e.g., transitory surface reflectance) and high-level non-natural properties (e.g., arbitrary category membership) typically reflect the causal relations of objects on shorter or longer temporal scales respectively. So, we should expect subpersonal representations of these features and properties to be less central to the explanations cognitive scientists offer for the agential and perceptual capacities that underlie perspectival unity. On the other hand, other properties are almost always going to be explanatorily relevant in these ways. Basic acoustic properties of sound convey a hight quantity of easily exploitable information about their sources and the wider environment, colours and shapes are highly predictive of object identity and behaviour on this privileged temporal scale (Marchi & Hohwy 2022). These facts are symptomatic of the cue-from-object dependencies that determine perceptual object category membership. There are distinctive reliable perceptual cues for, for example, ducks and trees, but not for categories constructed from instances of particular surface reflectance values or perspectival shapes – this is because only the former kind of object is associated with stable property appearances of this kind.

However, these generalisations are subject to exceptions. Idiosyncratic features of perceivers and occurrent features of particular perceptual situations can combine to give rise to causal

patterns on agentially relevant timescales from a much wider range of object properties and features. So, while object-appearance equivalence classes track perceptual object types, object appearance instances are not tied to a proprietary narrow range of basic perceptible properties. Easy solutions to the scope problem like those offered by supporters of intermediate-level processing theories can only answer questions about the types of perceptual object that are generally implicated in subpersonal explanations of phenomenal character. Answers to the scope question about particular experiences are only subject to the following much looser constraint:

> *Sensorimotor particularity*: An object, o, instances an appearance property A_p only if it has a property P that plays a causal role in the cue-from-object dependencies involved in a sensorimotor interaction Sϕo

Sensorimotor particularity makes explicit a consequence of *object-appearance property.* While it's trivially true that properties that standardly generate cue-from-object dependencies play the causal role of the corresponding perceptual object type, in a given perceptual situation another property might either generate or modify the same cue-supporting causal patterns in a perceiver's environment. Since interpretations are sensitive to contextual factors of this kind, there's no barrier to appearance properties that fail to generate perceptual object types that are not salient or useful in across-subject cognitive scientific studies.

While cognitive scientific evidence about perceptual mechanisms and personal-level perceptual capacities can put constraints on appearance properties, it cannot fix the properties and features that shape phenomenal character. But we've now seen that it can imposed principled constraints on it that are not derived from introspection. The question we should ask of candidate entries to the scope of perspectival experience is do they make a difference to the sensorimotor profile of the perceptual activity at hand. The sensorimotor acquaintance view is designed to explain and justify our attribution of singular content to people's perceptual experience and bases its account of phenomenal character in these sematic considerations. Since it's widely agreed that perceptual phenomenology has particular character[104] (corresponding to the singular content of perception and intentional state attributions central

[104] The particularity of perception has been a central theme of recent philosophical work on perceptual experience, both within the naïve realist framework (see, e.g., Gomes & French 2016) and outside of it (Siegel 2011; Schellenburg 2019, 2021). These debates are frequently conducted independently of explicit answers to the scope question (but see, Raleigh 2021). While some philosophers have questioned the claim that an account of perceptual experience must make space for particular content (see, e.g., Hill 2022), my arguments presuppose that the most important contest is over which account best accommodates particularity.

to perspectival experience) *sensorimotor particularity* offers a plausible constraint that orients the scope question towards the right target. The question is what falls withing the scope of particular (perspectival) experiences phenomenal character. The answer is those relational properties of perceptual objects that make a difference to the constitutive activity underlying corresponding experiences and which, when content matching these properties is appropriately attributed to them, qualify the experience as the accomplishment of a perspective on that object. While determining the bounds set by this constraint requires further, partially empirical, investigation it is already clear that they are less restrictive than alternatives derived from object appearance types tied to specific processing levels. They may turn out to be more restrictive than appearance types derived from introspection.

b) **The limits of sensorimotor acquaintance**

These considerations give rise to a limitation on sensorimotor acquaintance that the main contemporary proponents of the sensorimotor acquaintance view are unwilling to accept. Both Ward and Raleigh stress the intrinsic status of perceptible properties. For Raleigh perceived shape and colour are 'intrinsic physical features of objects' that are revealed to perceivers via more or less introspectively transparent modes of appearing,[105] while for Ward, "the scope and limits of your sensorimotor capacities constrain the set of objective properties you experience" (p. 268). In context, it's clear that Ward means by 'objective' what I've called the 'intrinsic' properties of an object. Towards the end of his discussion, he considers the objection that the experiences of sensorimotor perceivers are constituted only by "narcissistic properties" (Akins 1996), available to the subject in light of idiosyncratic features of her sensorimotor interactions with objects and not instanced across subjects. Avoiding this consequence is important if the view is going to offer something like the traditional naïve realist picture. Part of the motivation for naïve realism is a maximally non-narcissistic picture of the perceiving subject – someone unconstrained by Pater's "thick wall of personality" who delights in the greenness of the grass, there for all with eyes but to see it (see, above, §1.1a). But knowledge of intrinsic properties is not a feasible accomplishment for subjects equipped only with the resources that the

[105] At some points Raleigh moderates this claim. Talking about colour properties, for example, he notes that "in a perceptual experience of a coloured opaque surface we have awareness both of a non-relational feature of the physical surface and the relational property of how this surface feature looks" (p. S4097). He takes it that this allows rationalists to accommodate 'realist' and 'dispensationalist' views in the metaphysics of colour. I'm uncertain about this proposed accommodation – I think more needs to be said to justify mapping a personal-level distinctions between ways object look and the 'objective' colours they appear to have, to a more fundamental metaphysical distinction between an object's dispositions and intrinsic properties. I won't pursue the point, however, since the version of the sensorimotor acquaintance view involves denying that objects' intrinsic properties are perceived. If this is right, it subsumes the worry just expressed.

sensorimotor acquaintance account gives them. The reasons for this derive from the metaphysics of accomplishments and from the nature of the resources that sensorimotor perceivers bring to the task of getting acquainted with their surroundings. While knowledge of the intrinsic properties of things is a constitutive goal of perspectival experience, its attainment is at best a partial accomplishment. Some have suggested that considerations along these lines tell against the sensorimotor framework in general (see, e.g., Campbell 2008). The rest of this final section aims to make the opposite case.

First, let's consolidate the picture of the subject of sensorimotor acquaintance given above. Doing so will illustrate how the apparent problem arises and help to indicate why it is only apparent. The sensorimotor acquaintance view claims that acquaintance is realised by concrete, temporally extended interactions. This interaction is mediated by sensorimotor capacities. Capacities to enjoy experiences with the phenomenal character they have therefore depend on more fundamental capacities to maintain perceptual contact with environmental objects that partly constitute this character. These in turn depend partly on neurocomputational states of the perceiver (object files) but also on how the interactions thereby maintained actually play out. The two elements are interdependent, contents of object files affect the sensorimotor interactions which affect inputs, and so contents of object files – these dynamics are captured by characterising attentional and intentional states of the perceiver. While the acquaintance relation holds, the object plays a constitutive role in the subject's perceptual phenomenology because the object's appearing to the subject supervenes on the exercise of sensorimotor capacities (tracking the object via the coordination of sensory processing and motor skills). That is, it supervenes on a causal and informational feedback loop between subject and object, rather than on the intrinsic properties of either the object itself or a state subject, distally caused by the object.

This describes a causal structure common to 'good' cases – in which a subject encounters an ordinary environment – and 'bad' cases – in which casual factors in non-standard environments conspire to produce perceptual objects whose intrinsic properties radically differ from those of paradigms of the object-appearance class to which they belong. While objects are constitutive of the phenomenal character of perspectival experiences, they are so in virtue of the causal structure of subject-object interactions. This causal structure is not proprietary to the objects that interpretation uses to construct paradigms for experience. As Campbell rightly notes, the "content of the experience of seeing calls out to be characterized in terms of the objective, mind-independent, categorical properties of objects" (Campbell 2010, p. 199). But we should not conclude from this that the intrinsic properties that are ultimately responsible for the way

objects shape experiences of them are the same as the properties that show up in experience. Although Campbell worries that "[m]erely exhibiting perceptual constancy in perceptual representation [i.e. property attribution] is not enough for a grasp of the objective, intrinsic, categorical properties of the objects around one" (ibid.), the choice is not between experiential access either to intrinsic or 'narcissistic' properties. The relational properties that constitute perspectival experience are repeatable across experiences of multiple subjects. Since they are determined in part by the capacities exercised in grasping them, we should expect similarities in these capacities to correspond to phenomenal similarities. In the same way that we type objects of experience, we can type experiences.

The particularity of experience entails that at the most precise level of description, no two perspectival experiences will be exactly alike. Each is an unrepeatable episode in the psychological history of an individual subject. Unlike cognitive attitudes, they can be individuated by phenomenal properties that reflect the contextual and historical nature of their subpersonal constituents – episodes of representation-guided sensorimotor interaction with environmental particulars. Neither this phenomenal character, nor introspection on it constitutes knowledge of the intrinsic properties of the objects of experience. Such knowledge is accessible only from a 'sideways on' (McDowell 1994) view of the accomplishment of getting acquainted. If this is something available to subjects of acquaintance, it requires further conditions to be met. The way objects show up in perspectival experience is necessarily partial and context-bound. Their appearances do not disclose their intrinsic properties, but in specifying them via our interpretive grasp on the structure of interaction mediated experience we reveal objective, interpersonally accessible patterns that make up the subjective perspectives on the world inhabited by perceivers like us.

4) Conclusion

Does the resulting account deserve to be called a form of naïve realism? Does the sensorimotor acquaintance view represent a new (or, maybe, emerging) way to be a naïve realist about perceptual experience? I don't think so. This isn't just because sensorimotor acquaintance theorists can't adopt the full suite of intuitions that have traditionally motivates naïve realism proper – though this certainly should play a role. It's also because there are some highly non-commonsensical elements to the view as it's been articulated here, and it's far from clear that these can be eliminated while retaining its distinctive explanatory payoff. To think of the structure of everyday experience as determined by a complex pattern of causal unity between neurocomputational mechanisms and environmental structures sensorimotor agency, even if

our explanatory labels for these interacting elements derive from the vocabulary of ordinary psychological interpretation and explanation, is an insight (if it is one at all) that's not in the gift of naïve intuitions about perception, action and experience. Even if this picture situates us, as I've argued it does, out in the open among the people, happenings and things that we pretheoretically seem to share the world with, it has as much in common with the Kantian and neo-Kantian models of the experiencing subject that early acquaintance-theorists reacted against as it has with the common sense and ordinary language accounts that they helped to foster. Despite recent attempts to cast Kant as a naïve realist (Gomes 2014, 2017; Stephenson 2015; Allen 2020), it's hard to occupy these two positions comfortably and at the same time. In the last section, I've tried to raise some doubts about our motivation for doing so.

A popular rhetorical move at least since Berkeley, who presented his radically counterintuitive and arguably anti-realist account (although, see Pappas 1991) as an attempt to "unite, and place in a clearer light, that truth which was before shared between the vulgar and the philosophers" (2012), has been for philosophers to themselves as the defenders of the obvious, everyday and manifest against the bad faith of speculative philosophy. But these advertisements typically present the public-facing side of a lot of intricate theorising – if not downright speculation – behind the scenes. And in the case of naïve realism, this behind-the-scenes work has frequently involved serious sympathetic engagement with scientific efforts to go beyond the ordinary and obvious. Like Berkeley's *Essay*, Reid's *Inquiry* contains a sustained discussion of the contemporary state of the art in optics, and more recent naïve realists have defended the account in detailed cognitive- and neuroscientific terms. While this attests to the sophistication of naïve realism's defenders, I think it should lead us to be skeptical of the alleged opposition that philosophers' praise of naïveté implicitly or explicitly invokes. If my argument has been on track, acquaintance is an everyday accomplishment, but accounting for it requires an interlevel explanatory framework that builds upon the naïve perspective. One of the levels might track common sense, but this common sense is not secure. Interdependent with underlying mechanistic explanations, it will periodically be subject to disruption from below.

Maybe this is a case in which to save the appearances in the face of skeptical (scientific) and idealist (philosophical) challenges we must go beyond them. But it's worth remembering that going beyond appearances we might not leave them all intact or arranged as we had taken them to be. Anscombe (1959, p. 151) recounts a conversation with Wittgenstein:

> He once greeted me with the question: 'Why do people say that it was natural to think the sun went around the earth rather than that the earth turned on its axis?'

> I replied: 'I suppose because it looks as if the sun went around the earth.' 'Well', he asked 'what would it have looked like if it looked as if the earth turned on its axis?' My reply was to hold my hands with the palms upward and raise them with a circular sweep at the same time leaning backwards and assuming a dizzy expression. 'Exactly!' he said.

I take it that the point here is that philosophy leaves everything where it is – that the appearances were in good order before bad faith philosophy got hold of them. But don't we, by gaining the shift in perspective that the Copernican revolution enacted see the all the same phenomena in a new light? This is the question I take from the famous parallel metaphor Kant uses to begin his first *Critique.* For Anscombe, the notion that there would be a difference between the two *objective* ways the world would look is "really thoughtlessness" that could be "destroyed by a [i.e., Wittgenstein's] question" (ibid.). This thesis has attempted to articulate a way in which these two questions bear on each other. Subjects bring their own capacities and idiosyncrasies to their acquaintance with objects, and this means that what emerges in experience can only be a view of the object from within the limitations of a particular perspective – one that imposes a requirement for what Langton (1998) has called Kant's 'epistemic humility'. I've argued that the sensorimotor acquaintance view gives us a picture of what phenomenal acquaintance, conditioned by this kind of humility, could be.

Bibliography

Adams, F. (2010). Action theory meets embodied cognition. *Causing human actions: New perspectives on the causal theory of action*, 229-52.

Aguilar, J. H., & Buckareff, A. A. (2010*). Causing human actions: New perspectives on the causal theory of action*. MIT press

Aizawa, K. (2007). Understanding the embodiment of perception. *The Journal of philosophy*, *104*(1), 5-25.

Akins, K. (1996) 'Of Sensory Systems and the "Aboutness" of Mental States', The Journal of Philosophy, 93/7: 337–72

Allen, K. (2020). The value of perception. *Philosophy and Phenomenological Research*, *100*(3), 633-656.

Allen, M., & Friston, K. J. (2018). From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese*, *195*(6), 2459-2482.

Allport, A. (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years.

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological science*, *15*(2), 106-111.

Alvarez, M., & Hyman, J. (1998). Agents and their Actions. *Philosophy*, *73*(2), 219-245.

Amaya, S. (2017). Basic actions reloaded. *Philosophy Compass*, *12*(9), e12435.

Anscombe, G. E. M. (1957). *Intention*. Blackwell.

Anscombe, G. E. M. (1959). An introduction to Wittgenstein's Tractatus. Harper & Row

Anscombe, G. E. M. (1965). The intentionality of sensation: A grammatical feature. *Vision and mind: Selected readings in the philosophy of perception*, 55-75.

Arnott, S. R., Cant, J. S., Dutton, G. N., & Goodale, M. A. (2008). Crinkling and crumpling: an auditory fMRI study of material properties. *Neuroimage*, *43*(2), 368-378.

Arstila, V. (2018). Temporal experiences without the specious present. *Australasian Journal of Philosophy*, *96*(2), 287-302.

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological science*, *18*(7), 622-628.

Ayer, A. J. (1940) *Foundations of Empirical Knowledge*. MacMillan

Ásta. (2017). Social Kinds. In *The Routledge Handbook of Collective Intentionality* (pp. 290-299). Routledge.

Bacharach, J. (2021). Are Events Things of the Past?. *Mind*, *130*(518), 381-412.

Baker, L. R. (1997). Why constitution is not identity. *The Journal of Philosophy*, *94*(12), 599-621.

Ball, B. (2022). Intentionality, Point of View, and the Role of the Interpreter. *Phenomenology and Mind*, (22), 92-102.

Batty, C. (2010). Olfactory experience II: Objects and properties. *Philosophy Compass*, *5*(12), 1147-1156.

Bechtel, W. (1985). Realism, instrumentalism, and the intentional stance. *Cognitive Science*, *9*(4), 473-497.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 421-441.

Beck, J. (2015). Analogue magnitude representations: A philosophical introduction. *The British Journal for the Philosophy of Science*.

Beck, O. (2019). Rethinking naive realism. *Philosophical Studies*, *176*(3), 607-633.

Behrmann, M., Zemel, R. S., & Mozer, M. C. (1998). Object-based attention and occlusion: evidence from normal participants and a computational model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(4), 1011.

Bennett, M. R., Hacker, P. M. S., & Bennett, M. R. (2003). *Philosophical foundations of neuroscience* (Vol. 79). Blackwell.

Berkeley, G. (2012). *Three dialogues between Hylas and Philonous*. Broadview Press.

Bermúdez, J. L. (2000). Personal and sub-personal; A difference without a distinction. *Philosophical Explorations*, *3*(1), 63-82.

Bishop, J. (1983). Agent-causation. *Mind*, *92*(365), 61-79.

Bizley, J. K., & Cohen, Y. E. (2013). The what, where and how of auditory-object perception. *Nature Reviews Neuroscience*, *14*(10), 693-707.

Block, N. (2012). Discussion of J. Kevin O'Regan's "Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness". *Review of Philosophy and Psychology*, *3*(1), 89-108.

Brefczynski-Lewis, J. A., & Lewis, J. W. (2017). Auditory object perception: a neurobiological model and prospective review. *Neuropsychologia*, *105*, 223-242.

Bregman, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. MIT press.

Bricke, J. (1985). Locke, Hume and the Nature of Volitions. *Hume Studies*, *1985*(1), 15-51.

Broad, C. D. (1952) Some Elementary Reflections on Sense-Perception. *Philosophy* 27(1), 3-17.

Broadbent, D (1958). *Perception and communication*. London: Pergamon Press

Brogaard, B., & Chomanski, B. (2015). Cognitive penetrability and high-level properties in perception: unrelated phenomena?. *Pacific Philosophical Quarterly*, *96*(4), 469-486.

Brook, A. (2004). Kant, cognitive science and contemporary neo-kantianism. *Journal of Consciousness Studies*, *11*(10-11), 1-25.

Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, *195*(6), 2417-2444.

Buehler, D. (2019). Flexible occurrent control. *Philosophical Studies*, *176*(8), 2119-2137.

Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, *33*(1), 1-78.

Burge, T. (2009). Perceptual objectivity. *Philosophical Review*, *118*(3), 285-324.

Burge, T. (2011). Disjunctivism again. *Philosophical Explorations*, *14*(1), 43-80.

Burge, T. (2022). *Perception: First form of mind*. Oxford University Press.

Burr, D., & Thompson, P. (2011). Motion psychophysics: 1985–2010. *Vision research*, *51*(13), 1431-1456.

Byrne, A., & Siegel, S. (2017). Rich or thin?. In B Nanay (ed.). *Current controversies in philosophy of perception*, 59-80.

Call, J. (2009). Contrasting the social cognition of humans and nonhuman apes: the shared intentionality hypothesis. *Topics in Cognitive Science*, *1*(2), 368-379.

Campbell, J. (2002). *Reference and consciousness*. Clarendon Press.

Campbell, J. (2008). Sensorimotor knowledge and naïve realism. *Philosophy and Phenomenological Research*, *76*(3), 666-673.

Campbell, J. (2009). Consciousness and reference. In (Eds) McLaughlin, Beckermann & Walter *The Oxford Handbook of Philosophy of Mind*. OUP

Campbell, J. (2010). Relational View of Experience, and the Proximality Principle. *New essays on singular thought*, 193.

Campbell, J. (2011). Origins of Objectivity. *The Journal of Philosophy*, *108*(5), 269-285.

Campbell, J. (2020). Does that which makes the sensation of blue a mental fact escape us?. In *The Routledge handbook of philosophy of colour* (pp. 408-420). Routledge.

Campbell, J., & Cassam, Q., (2014). Berkeley's Puzzle. OUP

Camus, T., Hommel, B., Brunel, L., & Brouillet, T. (2018). From anticipation to integration: the role of integrated action-effects in building sensorimotor contingencies. *Psychonomic bulletin & review*, *25*, 1059-1065.

Carey, S., & Xu, F. (2001). Infants' knowledge of objects: Beyond object files and object tracking. *Cognition*, *80*(1-2), 179-213.

Carlyon, R. P. (2004). How the brain separates sounds. *Trends in cognitive sciences*, *8*(10), 465-471.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision research*, *51*(13), 1484-1525.

Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1993). A critique of pure vision.

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1585-1599.

Clark, A. (1993). *Sensory Qualities*. OUP

Clark, A. (2006). Attention & inscrutability: A commentary on John Campbell, reference and consciousness for the Pacific APA Meeting, Pasadena, California, 2004. *Philosophical Studies*, *127*(2), 167-193.

Clark, A. (2015a). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Clowes, R. W., & Chrisley, R. (2012). Virtualist representation. *International Journal of Machine Consciousness*, *4*(02), 503-522.

Cohen, J. (2015). Perceptual constancy. *Oxford Handbook of the Philosophy of Perception*

Colombo, M. (2013). Constitutive relevance and the personal/subpersonal distinction. *Philosophical Psychology*, *26*(4), 547-570.

Corrado, G. S., Sugrue, L. P., Brown, J. R., & Newsome, W. T. (2009). The trouble with choice: studying decision variables in the brain. In *Neuroeconomics* (pp. 463-480). Academic Press.

Crane, T. (Ed.). (1992). *The contents of experience: Essays on perception*. Cambridge University Press.

Craver, C. F., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology & philosophy*, *22*, 547-563.

Crowther, T. (2009). Watching, sight, and the temporal shape of perceptual activity. *Philosophical Review*, *118*(1), 1-27.

Cummins, R. (2000). " How does it work?" vs." What are the laws?" Two conceptions of psychological explanation. In F. Keil and R. Wilson (eds), *Explanation and Cognition*, MIT Press, 2000, pp 117-145.

Dainton, B. (2006). *Stream of consciousness: Unity and continuity in conscious experience*. Taylor & Francis.

Davidson, D. (1980). *Essays on Actions and Events.* Oxford University Press.

Davidson, D. (1984) *Inquires into Truth and Interpretation*. Oxford University Press

de Haan, E. H., Corballis, P. M., Hillyard, S. A., Marzi, C. A., Seth, A., Lamme, V. A., ... & Pinto, Y. (2020). Split-brain: what we know now and why this is important for understanding consciousness. *Neuropsychology review*, *30*(2), 224.

Dennett, D. C. (1968). *Content and Consciousness*. Routledge

Dennett, D. C. (1978). *Brainstorms: Philosophical essays on mind and psychology*. MIT press.

Dennett, D. C. (1987). *The intentional stance*. MIT press.

Dennett, D. C. (2000). Introduction. In G. Ryle. *The concept of mind*. Penguin.

Dennett, D. C. (2003). Who's on first? Heterophenomenology explained. Journal of Consciousness Studies, 10(9-10), 19-30.

Dennett, D. C. (2009). The part of cognitive science that is philosophy. *Topics in Cognitive Science*, *1*(2), 231-236.

Dennett, D. C. (2010). *Content and consciousness*. Routledge

Dennett, D. C. (2010). *Content and consciousness*. Routledge.n.

Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences*, *15*(2), 183-201

Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *353*(1373), 1245-1255.

Dewhurst, J. (2021). Folk psychological and neurocognitive ontologies. *Neural mechanisms: New challenges in the philosophy of neuroscience*, 311-334.

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the cognitive sciences*, *4*(4), 429-452.

Dickie, I. (2010). We are acquainted with ordinary things. *New essays on singular thought*.

Dickie, I. (2015). *Fixing reference*. Oxford University Press.

Dickie, I. (2016). The essential connection between epistemology and the theory of reference. *Philosophical Issues*, *26*, 99-129.

Ding, N., & Simon, J. Z. (2012). Emergence of neural encoding of auditory objects while listening to competing speakers. *Proceedings of the National Academy of Sciences*, *109*(29), 11854-11859.

Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical perspectives*, *26*, 1-18.

Drayson, Z. (2018). Extended minds and prime mental conditions: probing the parallels.

Duncan, M. (2021). Acquaintance. *Philosophy Compass*, *16*(3), e12727.

Dupré, J. (2017). Natural kinds. *A Companion to the Philosophy of Science*, 311-319.

Eadie, M. J. (2009). Hubert Airy, contemporary men of science and the migraine aura. *The Journal of the Royal College of Physicians of Edinburgh*, *39*(3), 263-267.

Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection*. Basic books.

Engel, A. K., Friston, K. J., & Kragic, D. (Eds.). (2015). *The pragmatic turn: Toward action-oriented views in cognitive science* (Vol. 18). MIT Press.

Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in cognitive sciences*, *17*(5), 202-209.

Epstein, P. F. (2022). In defense of the armchair: Against empirical arguments in the philosophy of perception. *Noûs*.

Evans, G. (1982*) The Varieties of Reference*. OUP

Fang, F., Kersten, D., & Murray, S. O. (2008). Perceptual grouping and inverse fMRI activity patterns in human visual cortex. *Journal of vision*, *8*(7), 2-2.

Feigenson, L., Carey, S., & Hauser, M. (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological science*, *13*(2), 150-156.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and brain sciences*, *39*.

Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Brooks/Cole.

Fodor, J. A. (1968). The appeal to tacit knowledge in psychological explanation. *The Journal of Philosophy*, *65*(20), 627-640.

Fodor, J. A. (1975). *The language of thought* (Vol. 5). Harvard university press.

Fodor, J., & Pylyshyn, Z. (1981). How direct is visual perception?: Some reflections on Gibson's "ecological approach." *Cognition,9*(2), 139–196.

Fothergill, John. "Remarks on That Complaint Commonly Known as Sick-Headache. Read Dec. 14, 1778." Medical Observations and Inquiries 6 (1757–1784): 103–137

Foxhall, K. (2014). Making modern migraine medieval: men of science, Hildegard of Bingen and the life of a retrospective diagnosis. *Medical History*, *58*(3), 354-374.

Foxhall, K. (2019). *Migraine: A History*. Johns Hopkins University Press.

Frankish, K. (2006). Review of Consciousness in Action, by Susan Hurley. *Mind*, *115*.

French, C., & Gomes, A. (2019). How naïve realism can explain both the particularity and the generality of experience. The Philosophical Quarterly, 69(274), 41-63.

Fridland, E. (2017). Automatically minded. *Synthese*, *194*(11), 4337-4363.

Friston, K., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. *Frontiers in psychology*, *3*, 151.

Gao, T., & Scholl, B. J. (2010). Are objects required for object-files? Roles of segmentation and spatiotemporal continuity in computing object persistence. *Visual Cognition*, *18*(1), 82-109.

Gazzaniga, M. S. (2005). Forty-five years of split-brain research and still going strong. *Nature Reviews Neuroscience*, *6*(8), 653-659.

Genone, J. (2014). Appearance and illusion. *Mind*, *123*(490), 339-376.

Genone, J. (2016). Recent work on naive realism. *American Philosophical Quarterly*, 1-25.

Gerrans, P. (2004). Cognitive architecture and the limits of interpretationism. *Philosophy, Psychiatry, & Psychology*, *11*(1), 43-48.

Gläscher, J. P., & O'Doherty, J. P. (2010). Model-based approaches to neuroimaging: combining reinforcement learning theory with fMRI data. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(4), 501-510.

Goadsby, P. J., Holland, P. R., Martins-Oliveira, M., Hoffmann, J., Schankin, C., & Akerman, S. (2017). Pathophysiology of migraine: a disorder of sensory processing. *Physiological reviews*.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Causal Theories of Mind*, 174.

Gomes, A. (2014). Kant on perception: Naive realism, non-conceptualism, and the B-deduction. The Philosophical Quarterly, 64(254), 1-19.

Gomes, A. (2016). Unity, objectivity, and the passivity of experience. European Journal of Philosophy, 24(4), 946-969.

Gomes, A. (2016). Unity, objectivity, and the passivity of experience. *European Journal of Philosophy*, *24*(4), 946-969.

Gomes, A. (2017). Naïve realism in Kantian phrase. Mind, 126(502), 529-578.

Gomes, A. (2021). Unity and objectivity in Strawson and Cassam. Analytic Philosophy, 62(1), 84-96.

Gomes, A., & French, C. (2016). On the particularity of experience. Philosophical Studies, 173(2), 451-460.

Goodale, M. (1997). Pointing the way to a unified theory of action and perception. *Behavioral and Brain Sciences*, *20*(4), 749-750.

Gray, C. M., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences*, *86*(5), 1698-1702.

Green, E. J. (2018). What do object files pick out?. *Philosophy of Science*, *85*(2), 177-200.

Green, E. J. (2019). A theory of perceptual objects. *Philosophy and Phenomenological Research*, *99*(3), 663-693.

Green, E. J., & Quilty-Dunn, J. (2021). What is an object file?. *The British Journal for the Philosophy of Science*.

Gregory, D. (2013) *Showing, Sensing, and Seeming: Distinctively Sensory Representations and their Contents*. OUP

Griffin, N. (2012). Russell and Leibniz on the Classification of Propositions. In *New Essays on Leibniz Reception: In Science and Philosophy of Science 1800-2000* (pp. 85-127). Basel: Springer Basel.

Griffiths, T. D., & Warren, J. D. (2004). What is an auditory object?. *Nature Reviews Neuroscience*, *5*(11), 887-892.

Hacking, I. (1991). A tradition of natural kinds. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *61*(1/2), 109-126.

Hadjikhani, N., & Vincent, M. (2019). Neuroimaging clues of migraine aura. *The Journal of Headache and Pain*, *20*(1), 1-7.

Hadjikhani, N., & Vincent, M. (2021). Can you have a migraine aura without knowing it?. *Current Opinion in Neurology*, *34*(3), 350-355.

Hafri, A., & Firestone, C. (2021). The perception of relations. *Trends in Cognitive Sciences*, *25*(6), 475-492.

Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience*, *18*(4), 196-207.

Hansen, C., & Rey, G. (2016). Files and Singular Thoughts Without Objects or Acquaintance: The Prospects of Recanati's (and Others')"Actualism". *Review of Philosophy and Psychology*, *7*, 421-436.

Hansen, J. M., Baca, S. M., VanValkenburgh, P., & Charles, A. (2013). Distinctive anatomical and physiological features of migraine aura revealed by 18 years of recording. *Brain*, *136*(12), 3589-3595.

Hayhoe, M. M. (2017). Vision and action. *Annual review of vision science*, *3*, 389-413.

Hesp, C., Ramstead, M., Constant, A., Badcock, P., Kirchhoff, M., & Friston, K. (2019). A multi-scale view of the emergent complexity of life: A free-energy proposal. In *Evolution, Development and Complexity* (pp. 195-227). Springer, Cham.

Hill, C. S. (2022). Susanna Schellenberg on perception. *Mind & Language*, *37*(2), 208-218.

Hillman, E. M. (2014). Coupling mechanism and significance of the BOLD signal: a status report. *Annual review of neuroscience*, *37*, 161.

Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, *36*(5), 791-804.

Hohwy, J. (2016). The self-evidencing brain. *Noûs*, *50*(2), 259-285.

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in cognitive sciences*, *8*(11), 494-500.

Hommel, B., Chapman, C. S., Cisek, P., Neyedli, H. F., Song, J. H., & Welsh, T. N. (2019). No one knows what attention is. *Attention, Perception, & Psychophysics*, *81*, 2288-2303.

Hornsby, J. (1980). *Actions*. Routledge

Hornsby, J. (2000). Personal and sub-personal; A defence of Dennett's early distinction. *Philosophical Explorations*, *3*(1), 6-24.

Hornsby, J. (2008). Agency and alienation.

Hornsby, J. (2010) The standard story of action: an exchange. In: Aguilar, J.H. and Buckareff, A.A. (eds.) *Causing Human Action: New Perspectives on the Causal Theory of Action.* MIT Press

Hornsby, J. (2013). Basic activity. In *Aristotelian Society Supplementary Volume* (Vol. 87, No. 1, pp. 1-18). Oxford University Press.

Hornsby, J. (2022). Naturalism in the Philosophy of Action. In M. Boyle & E. Mylonaki (Ed.), *Reason in Nature: New Essays on Themes from John McDowell* (pp. 171-190). Harvard University Press

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, *160*(1), 106.

Humphreys, G. W. (2016). Feature confirmation in object perception: Feature integration theory 26 years on from the Treisman Bartlett lecture. *The Quarterly Journal of Experimental Psychology*, *69*(10), 1910-1940.

Hurley, S. (1989) *Natural Reasons: Personality and Polity*. OUP

Hurley, S. (1998). *Consciousness in action*. Harvard University Press.

Hurley, S. (2003a). Action, the unity of consciousness, and vehicle externalism. *The unity of consciousness: Binding, integration, and dissociation*, 78-91.

Hurley, S. (2003b). Animal action in the space of reasons. *Mind & Language*, *18*(3), 231-257.

Hurley, S. (2008). The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and brain sciences*, *31*(1), 1-22.

Husserl, E. (1998) *Thing and Space: lectures of 1907.* Springer

Hutto, D. D. (2005). Knowing what? Radical versus conservative enactivism. *Phenomenology and the Cognitive Sciences*, *4*(4), 389-405.

Hylton, P. (2003). The theory of descriptions. *The Cambridge Companion to Bertrand Russell*, 202-40. Cambridge University Press.

Hyman, J. (2015). *Action, knowledge, and will*. Oxford University Press.

Inman, B. A. (1981). The intellectual context of Walter Pater's "conclusion". *Prose Studies*, *4*(1), 12-30.

Jackendoff, R. (1987). *Consciousness and the computational mind*. MIT Press.

James, W. (1890). *The Principles of Psychology*. Macmillan

Johnson, S. P. (2011). Development of visual perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(5), 515-528.

Jordan, K. E., Clark, K., & Mitroff, S. R. (2010). See an object, hear an object file: Object correspondence transcends sensory modality. *Visual Cognition*, *18*(4), 492-503.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, *24*(2), 175-219.

Kaiser, J., Bühler, M., & Lutzenberger, W. (2004). Magnetoencephalographic gamma-band responses to illusory triangles in humans. *Neuroimage*, *23*(2), 551-560.

Kalis, A. (2019). No intentions in the brain: A Wittgensteinian perspective on the science of intention. *Frontiers in psychology*, *10*, 946.

Kalderon, M. E. (2011). Before the law. *Philosophical issues*, *21*(1), 219-244.

Kant, I. (1998) Critique of Pure Reason. P. Guyer & A.W. Wood (trans.). Cambridge University Press.

Kaya, E. M., & Elhilali, M. (2017). Modelling auditory attention. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1714), 20160101.

Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (1999). Attention without awareness in blindsight. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *266*(1430), 1805-1811.

Kentridge, R. W., Nijboer, T. C., & Heywood, C. A. (2008). Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia*, *46*(3), 864-869.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.*, *55*, 271-304.

Kim, J. (1993). *Supervenience and mind: Selected philosophical essays*. Cambridge University Press.

Kirchhoff, M. D. (2018). Autopoiesis, free energy, and the life–mind continuity thesis. *Synthese*, *195*(6), 2519-2540.

Kirchhoff, M. D., & Hutto, D. D. (2016). Never mind the gap: Neurophenomenology, radical enactivism, and the hard problem of consciosuness.

Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended consciousness and predictive processing: A third wave view*. Routledge.

Knowles, J., & Raleigh, T. (Eds.). (2019). *Acquaintance: new essays*. Oxford University Press.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307-321.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307.

Korsgaard, C. M. (1989). Personal identity and the unity of agency: A Kantian response to Parfit. *Philosophy & Public Affairs*, 101-132.

Korsgaard, C. M. (2009). *Self-constitution: Agency, identity, and integrity*. OUP Oxford.

Kristjánsson, Á., & Egeth, H. (2020). How feature integration theory integrated cognitive psychology, neurophysiology, and psychophysics. *Attention, Perception, & Psychophysics*, *82*, 7-23.

Kubovy, M., & Van Valkenburg, D. (2001). Auditory and visual objects. *Cognition*, *80*(1-2), 97-126.

Lacey, S., Tal, N., Amedi, A., & Sathian, K. (2009). A putative model of multisensory object representation. *Brain topography*, *21*, 269-274.

Langton, R. (1998). *Kantian humility: Our ignorance of things in themselves*. Oxford University Press.

Lavin, D. (2013). Must there be basic action?. *Noûs*, *47*(2), 273-301.

Lavin, D. (2015). Action as a form of temporal unity: on Anscombe's Intention. *Canadian Journal of Philosophy*, *45*(5-6), 609-629.

Leao, A. A. P., & Morison, R. S. (1945). Propagation of spreading cortical depression. *Journal of Neurophysiology*, *8*(1), 33-45.

Lederman, S. J., & Klatzky, R. L. (1993). Extracting object properties through haptic exploration. *Acta psychologica*, *84*(1), 29-40.

Levine, J. (2010). Demonstrative thought. *Mind & Language*, *25*(2), 169-195.

Lewis, D. (1974a). Causation. *The journal of philosophy*, *70*(17), 556-567.

Lewis, D. (1974b). Radical interpretation. *Synthese*, 331-344.

Libet, B. (1999). Do we have free will?. *Journal of consciousness studies*, *6*(8-9), 47-57.

Locke, E. (2000). Motivation, cognition, and action: An analysis of studies of task goals and knowledge. *Applied Psychology*, *49*(3), 408-429.

Lycan, W. G. (1981). Toward a homuncular theory of believing. *Cognition and Brain Theory*, *4*(2).

Lycan, W. (1991). Homuncular Functionalism Meet PDP. In Ramsey, Stich and Rumelhart, (eds.), *Philosophy and Connectionist Theory*. L. Erlbaum.

Madary, M. (2017). *Visual phenomenology*. MIT Press.

Malpas, J. (1999). Constituting the mind: Kant, Davidson and the Unity of Consciousness. *International journal of philosophical studies*, *7*(1), 1-30.

Marchi, F., & Hohwy, J. (2022). The intermediate scope of consciousness in the predictive mind. *Erkenntnis*, *87*(2), 891-912.

Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. MIT Press.

Martin, M. (2002). The transparency of experience. *Mind & Language*, *17*(4), 376-425.

Martin, M. G. (2006). On being alienated. *Perceptual experience*, 354-410.

Martin, M. (2007). Austin: Sense & Sensibilia Revisited.

Martin, M. G. (1998). Setting things before the mind. *Royal Institute of Philosophy Supplements*, *43*, 157-179.

McCabe, D. P., & Castel, A. D. (2008). Seeing is believing: The effect of brain images on judgments of scientific reasoning. *Cognition*, *107*(1), 343-352.

McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quarterly (1950-)*, *44*(175), 190-205.

McDowell, J. (1996). *Mind and world*. Harvard University Press.

McDowell, J. (2009). Conceptual capacities in perception. *Having the world in view: Essays on Kant, Hegel, and Sellars*, *127*, 144. Harvard University Press.

Meltzoff, A. N. (2013). *Origins of social cognition. Navigating the social world*, 139-144.

Menary, R. (2016). Pragmatism and the pragmatic turn in cognitive science. *The pragmatic turn: Toward action-oriented views in cognitive science*, 215-234.

Merleau-Ponty, M. (2013). *Phenomenology of perception*. Routledge.

Metzinger, T. K. (2013). The myth of cognitive agency: subpersonal thinking as a cyclically recurring loss of mental autonomy. *Frontiers in psychology*, *4*, 931.

Metzinger, T., & Wiese, W. (2017). Philosophy and predictive processing. *Frankfurt am Main: MIND Group*.

Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT press.

Millikan, R. G. (1990). The myth of the essential indexical. *Noûs*, 723-734.

Millikan, R. G. (2004). *Varieties of meaning: the 2002 Jean Nicod lectures*. MIT press.
Milner, D., & Goodale, M. (2006). *The visual brain in action* (Vol. 27). OUP Oxford.
Moll, H., & Meltzoff, A. N. (2011). Foundation in Joint Attention. *Perception, causation, and objectivity*, *286*, 286-304.
Montero, B. G. (2016). *Thought in action: Expertise and the conscious mind*. Oxford University Press.
Moore, G. E. (1903). The refutation of idealism. *Mind*, *12*(48), 433-453.
Moore, G. E., & Stout, G. F. (1913, January). Symposium: The Status of Sense-Data. In *Proceedings of the Aristotelian Society* (Vol. 14, pp. 355-406). Aristotelian Society, Wiley.
Muñoz-Suárez, C., & De Brigard, F. (Eds.). (2015). *Content and Consciousness Revisited: With Replies by Daniel Dennett* (Vol. 7). Springer.
Muñoz-Suárez, C., & De Brigard, F. (Eds.). (2015). *Content and Consciousness Revisited: With Replies by Daniel Dennett* (Vol. 7). Springer.
Murray, M. (1973). Heidegger and Ryle: Two versions of phenomenology. *The Review of Metaphysics*, *27*(1), 88-111.
Nagel, T. (1974). What is it like to be a bat?. *The philosophical review*, *83*(4), 435-450.
Neta, R. (2008). In defence of disjunctivism. *Disjunctivism: Perception, action, knowledge*, 311-329.
Noë, A. (2004). *Action in perception*. MIT press.
Noë, A. (2008). Reply to Campbell, Martin, and Kelly. *Philosophy and Phenomenological Research*, *76*(3), 691-706.
Nudds, M. (2010). What are auditory objects?. *Review of Philosophy and Psychology*, *1*(1), 105-122.
O'Callaghan, C. (2012). Perception and multimodality. *Oxford handbook of philosophy of cognitive science*, 92-117.
O'Callaghan, C. (2015). The multisensory character of perception. *The Journal of Philosophy*, *112*(10), 551-569.
O'Hare, L., Asher, J. M., & Hibbard, P. B. (2021). Migraine Visual Aura and Cortical Spreading Depression—Linking Mathematical Models to Empirical Evidence. *Vision*, *5*(2), 30.
O'Regan, J. K., & Noë, A. (2001a). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, *24*(5), 939-973.
O'Regan, J. K., & Noë, A. (2001b). What it is like to see: A sensorimotor theory of perceptual experience. *Synthese*, *129*, 79-103.
O'Shea, J. R. (2021). What is the myth of the given?. Synthese, 199(3), 10543-10567.
Orlandi, N. (2016). Bayesian perception is ecological perception. *Philosophical Topics*, *44*(2), 327-352.
Pappas, G. S. (1991). Berkeley and common sense realism. *History of Philosophy Quarterly*, *8*(1), 27-42.
Parton, A., Malhotra, P., & Husain, M. (2004). Hemispatial neglect. *Journal of Neurology, Neurosurgery & Psychiatry*, *75*(1), 13-21.
Pater, W. (1873). *Studies in the History of the Renaissance*. (Ed.) A. Drake. E-Texts for Victorianists [2001] downloaded from http://ajdrake.com/wiki/victorianprose/pater_renaissance_1873.pdf 25/03/2023
Pautz, A. (2013). Do the benefits of naïve realism outweigh the costs? Comments on fish, perception, hallucination and illusion. *Philosophical Studies*, *163*(1), 25-36.
Pautz, A. (2014). The real trouble with phenomenal externalism: New empirical evidence for a brain-based theory of consciousness. *Consciousness inside and out: Phenomenology, neuroscience, and the nature of experience*, 237-298.
Pautz, A. (2021a) *Perception*. Taylor & Francis.
Pautz, A. (2021b). Varieties of Interpretationism about Belief and Desire. *Analysis*, *81*(3), 512-524.

Pautz, A. (2023). Naïve realism v representationalism: an argument from science. *Contemporary Debates in Philosophy of Mind. Hoboken: Blackwell*.

Pepp, J. (2020). Is Dickie's Account of Aboutness-Fixing Explanatory?. *Theoria*, *86*(6), 801-820.

Perry, J. (1979). The problem of the essential indexical. *Noûs*, 3-21.

Pezzulo, G., & Cisek, P. (2016). Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends in cognitive sciences*, *20*(6), 414-424.

Phillips, R. L. (1964). Austin and Berkeley on Perception. *Philosophy*, *39*(148), 161-163.

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data?. *Trends in cognitive sciences*, *10*(2), 59-63.

Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, *22*(3-4), 262-275.

Prinz, J. (2000). A neurofunctional theory of visual consciousness. *Consciousness and Cognition*, *9*(2), 243-259.

Prinz, J. (2013). Attention, atomism, and the disunity of consciousness. *Philosophy and Phenomenological Research*, *86*(1), 215-222.

Prinz, J. (2013). Siegel's get rich quick scheme. *Philosophical Studies*, *163*(3), 827-835.

Prinz, J. J. (2005). A neurofunctional theory of consciousness. *Cognition and the brain: Philosophy and neuroscience movement*, 381-96.

Prinz, J. J. (2005). A neurofunctional theory of consciousness. *Cognition and the brain: Philosophy and neuroscience movement*, 381-96.

Pylyshyn, Z. W. (2000). Situating vision in the world. *Trends in cognitive sciences*, *4*(5), 197-207.

Pylyshyn, Z. W. (2007). *Things and places: How the mind connects with the world*. MIT press.

Quine, W. V. O. (1960). *Word and object*. MIT press

Raleigh, T. (2019). The recent Renaissance of acquaintance. In J. Knowles, & T. Raleigh (Eds.). *Acquaintance: new essays*. Oxford University Press.

Raleigh, T. (2021). Visual acquaintance, action and the explanatory gap. *Synthese*, *198*(Suppl 17), 4081-4106.

Ramsey, W. M. (2007). *Representation reconsidered*. Cambridge University Press.

Ramstead, M. J., Kirchhoff, M. D., & Friston, K. J. (2020). A tale of two densities: Active inference is enactive inference. *Adaptive Behavior*, *28*(4), 225-239.

Rashbrook, O. (2013). Diachronic and synchronic unity. *Philosophical studies*, *164*(2), 465-484.

Ross, D. (2015). A most rare achievement: Dennett's scientific discovery in Content and Consciousness. *Content and Consciousness Revisited: With Replies by Daniel Dennett*, 29-48.

Ross, D., Sharp, C., Vuchinich, R. E., & Spurrett, D. (2012). *Midbrain mutiny: The picoeconomics and neuroeconomics of disordered gambling: Economic theory and cognitive science*. MIT press.

Rupert, R. D. (2018) The Self in the Age of Cognitive Science: Decoupling the Self from the Personal Level. *Exchange*, *47*(1), 2.

Russell, B. (1905). On denoting. *Mind*, *14*(56), 479-493.

Russell, B. (1910). Knowledge by acquaintance and knowledge by description. In *Proceedings of the Aristotelian society* (Vol. 11, pp. 108-128). Aristotelian Society, Wiley.

Russell, B. (1911). On the relations of universals and particulars. In *Proceedings of the Aristotelian society* (Vol. 12, pp. 1-24). Aristotelian Society, Wiley.

Russell, . B. (1912) *The Problems of Philosophy*. Williams and Nordgate

Russell, B. (1946). *History of Western Philosophy*. George Allen & Unwin.

Russell, B. (1958). My Philosophical Development. *Hibbert Journal*, *57*(a).

Sacks, O. (1987). *Migraine*. Picador.

Schechter, E. (2013). The unity of consciousness: Subjects and objectivity. *Philosophical Studies*, 165(2), 671-692.

Scheerer, E. (1994) Psychoneural isomorphism: Historical background and current relevance, Philosophical Psychology, 7:2, 183-210
Schellenberg, S. (2017). The origins of perceptual knowledge. *Episteme*, *14*(3), 311-328.
Schneegans, S., & Bays, P. M. (2019). New perspectives on binding in visual working memory. *British Journal of Psychology*, *110*(2), 207-244.
Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*(1-2), 1-46.
Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive psychology*, *38*(2), 259-290.
Schwitzgebel, E. (2002). A phenomenal, dispositional account of belief. *Noûs*, *36*(2), 249-275.
Schwitzgebel, E. (2006) Belief, *The Stanford Encyclopedia of Philosophy,* Edward N. Zalta (ed.)
Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review*, *117*(2), 245-273.
Sellars, W. (1956) *Empiricism and The Philosophy of Mind.* In H. Feigl and M. Scriven (eds.) *The Foundations of Science and the Concepts of Psychoanalysis, Minnesota Studies in the Philosophy of Science, Vol. I.* University of Minnesota Press
Sergent, C. (2018). The offline stream of conscious representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1755), 20170349.
Seth, A. K. (2015). Presence, objecthood, and the phenomenology of predictive perception. *Cognitive neuroscience*, *6*(2-3), 111-117.
Shagrir, O., & Bechtel, W. (2017). Marr's computational level and delineating phenomena. *Explanation and integration in mind and brain science*, 190-214.
Shamma, S. A., Elhilali, M., & Micheyl, C. (2011). Temporal coherence and attention in auditory scene analysis. *Trends in neurosciences*, *34*(3), 114-123.
Shinn-Cunningham, B. G. (2008). Object-based auditory and visual attention. *Trends in cognitive sciences*, *12*(5), 182-186.
Shinn-Cunningham, B., Best, V., & Lee, A. K. (2017). Auditory object formation and selection. *The auditory system at the cocktail party*, 7-40.
Siegel, S. (2011). *The contents of visual experience*. Oxford University Press.
Siegel, S. (2016). *The rationality of perception*. Oxford University Press.
Siegel, S. (2019). The uneasy heirs of acquaintance. *Philosophical Issues*, *29*(1), 348-365.
Silvanto, J. (2013). Transcranial magnetic stimulation and vision. *Handbook of Clinical Neurology*, *116*, 655-669.
Silverman, D. (2018). Bodily skill and internal representation in sensorimotor perception. *Phenomenology and the cognitive sciences*, *17*(1), 157-173.
Skinner, B. F. (1957). The experimental analysis of behavior. *American scientist*, *45*(4), 343-371.
Small, W. (2016). Bodily movement and its significance. *Philosophical Topics*, *44*(1), 183-206.
Smith, M. (2010) The standard story of action: an exchange. In: Aguilar, J.H. and Buckareff, A.A. (eds.) *Causing Human Action: New Perspectives on the Causal Theory of Action.*MIT Press
Smith, M. (2012). Four objections to the standard story of action (and four replies). *Philosophical Issues*, *22*, 387-401.
Snowdon, P. F. (2015). Sense-data. *The Oxford handbook of philosophy of perception*, 118-135.
Soteriou, M. (2009). The disjunctive theory of perception. Stanford Encyclopaedia of Philosophy
Soteriou, M. (2013). *The mind's construction: The ontology of mind and mental action*. Oxford University Press.
Spelke, E. S. (1990). Principles of object perception. *Cognitive science*, *14*(1), 29-56.
Spelke, E. S. (1990). Principles of object perception. *Cognitive science*, *14*(1), 29-56.
Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental science*, *10*(1), 89-96.
Spelke, E. S., Phillips, A., & Woodward, A. L. (1995). Infants' knowledge of object motion and human action.

Stephenson, A. (2015). Kant on the object-dependence of intuition and hallucination. *The Philosophical Quarterly*, *65*(260), 486-508.

Stevens, S. S. (1958). Problems and methods of psychophysics. *Psychological Bulletin*, *55*(4), 177.

Steward, H. (1997). *The Ontology of Mind: Events, Processes, and States*. Oxford University Press.

Steward, H. (2000). Do actions occur inside the body?. *Mind & Society*, *1*, 107-125.

Stich, S. P. (1983). *From folk psychology to cognitive science: The case against belief*. MIT press.

Strawson, P. F. (1950). On referring. *Mind*, *59*(235), 320-344.

Strawson, P. F. (1959) *Individuals: An Essay in Descriptive Metaphysics*. Methuen

Strawson, P. F. (1966) *The Bounds of Sense: An Essay on Kant's Critique of Pure Reason*. Methuen.

Sussman, E., Ritter, W., & Vaughan, H. G. (1999). An investigation of the auditory streaming effect using event-related brain potentials. *Psychophysiology*, *36*(1), 22-34.

Tallon-Baudry, C., & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in cognitive sciences*, *3*(4), 151-162.

Thomasson, A. L. (2002). Phenomenology and the development of analytic philosophy. *The Southern journal of philosophy*, *40*(Supplement), 115-142.

Thompson E (2007) Mind in life: biology, phenomenology, and the sciences of mind. Harvard University Press

Toner, J., Montero, B. G., & Moran, A. (2015). The perils of automaticity. *Review of General Psychology*, *19*(4), 431-442.

Toner, J., Montero, B. G., & Moran, A. (2016). Reflective and prereflective bodily awareness in skilled action. *Psychology of consciousness: theory, research, and practice*, *3*(4), 303.

Tóth, B., Kocsis, Z., Háden, G. P., Szerafin, Á., Shinn-Cunningham, B. G., & Winkler, I. (2016). EEG signatures accompanying auditory figure-ground segregation. *Neuroimage*, *141*, 108-119.

Kalderon, M. E., & Travis, C. (2013). Oxford realism. In Michael Beaney (ed.), Oxford Handbook of the History of Analytic Philosophy. Oxford University Press.

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological review*, *76*(3), 282.

Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. Quarterly Journal of Experimental Psychology, 40A, 201–237.

Treisman, A. (1996). The binding problem. *Current opinion in neurobiology*, *6*(2), 171-178.

Treisman, A.., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, *12*(1), 97-136.

Trenholm, S., & Krishnaswamy, A. (2020). An annotated journey through modern visual neuroscience. *Journal of Neuroscience*, *40*(1), 44-53.

Tye, M. (199f5). *Ten problems of consciousness: A representational theory of the phenomenal mind*. mit Press.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT press.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT press.

Velleman, J. D. (1992). What happens when someone acts?. *Mind*, *101*(403), 461-481.

Vendler, Z. (1957). Verbs and times. *The philosophical review*, *66*(2), 143-160.

Vrahimis, A. (2013). " Was There a Sun Before Men Existed?": AJ Ayer and French Philosophy in the fifties. *Journal for the History of Analytical Philosophy*, *1*(9).

Ward, D. (2016a). Achieving transparency: An argument for enactivism. Philosophy and Phenomenological Research, 93(3), 650-680.

Ward, D. (2016b). Hurley's transcendental enactivism. Journal of Consciousness Studies, 23(5-6), 12-38.

Ward, D. (2023). Sensorimotor Relationalism and Conscious Vision. The Philosophical Quarterly.
Ward, D., Silverman, D., & Villalobos, M. (2017). Introduction: The varieties of enactivism. *Topoi*, *36*(3), 365-375.
Watzl, S. (2011). The nature of attention. *Philosophy compass*, *6*(11), 842-853.
Watzl, S. (2023). What attention is. The priority structure account. *Wiley Interdisciplinary Reviews: Cognitive Science*, *14*(1), e1632.
Weatherall, M. W. (2021). From "Transient Hemiopsia" to Migraine Aura. *Vision*, *5*(4), 54.
Williams, J. R. G. (2020). *The metaphysics of representation*. Oxford University Press
Winkler, I., Denham, S. L., & Nelken, I. (2009). Modeling the auditory scene: predictive regularity representations and perceptual objects. *Trends in cognitive sciences*, *13*(12), 532-540.
Wittgenstein (1953) Philosophical Investigations. G. E. A. Anscombe (trans). Blackwell.
Wolter, L. (2009). Demonstratives in philosophy and linguistics. *Philosophy Compass*, *4*(3), 451-468.
Wong, H. Y. (2014). Personal and sub-personal: Overcoming explanatory apartheid. In *Communicative action* (pp. 93-104). Springer, Singapore.
Wu, W. (2011). Attention as selection for action. *Attention: Philosophical and psychological essays*, *97*.
Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, *440*(7080), 91-95.
Ylikoski, P. (2013). Causal and constitutive explanation compared. *Erkenntnis*, *78*, 277-297.
Zmigrod, S., & Hommel, B. (2009). Auditory event files: Integrating auditory perception and action planning. *Attention, Perception, & Psychophysics*, *71*(2), 352-362.
Zmigrod, S., & Hommel, B. (2013). Feature integration across multimodal perception and action: a review. *Multisensory research*, *26*(1-2), 143-157.