

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Space and Self-Awareness

Permalink

<https://escholarship.org/uc/item/2xr655cn>

Author

Schwenkler, John Louis

Publication Date

2009

Peer reviewed|Thesis/dissertation

Space and Self-Awareness

by

John Louis Schwenkler

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Alva Noë, Chair

Professor John Campbell

Professor Lynn Robertson

Professor Stephen Palmer

Fall 2009

Space and Self-Awareness

© 2009

by John Louis Schwenkler

Abstract

Space and Self-Awareness

by

John Louis Schwenkler

Doctor of Philosophy in Philosophy

University of California, Berkeley

Professor Alva Noë, Chair

How should we think about the role of visual spatial awareness in perception and perceptual knowledge? A common view, which finds a characteristic expression in Kant but has an intellectual heritage reaching back farther than that, is that an account of spatial awareness is fundamental to a theory of experience because spatiality is the defining characteristic of “outer sense”, of our perceptual awareness of how things are in the parts of the world that surround us. A natural counterpart to this idea is to treat self-consciousness as residing in a kind of sense that is fundamentally “inner”, such as introspection or whatever else gives one privileged access to his own mental states as well as the proprioceptive and kinesthetic awareness of bodily position. This division is compatible, of course, with the idea that inner sense provides an awareness of a distinctive kind of “body space”, but it treats that as importantly different from the awareness of the worldly space around one.

In contrast to such a picture, this dissertation proposes an account of visual spatial awareness according to which it is no less a source of self-consciousness than of the awareness of the objects around us, and an account of self-awareness in which visual experience is essentially implicated. I begin by arguing that we should think of visual spatial awareness not as necessary for the individuation of visual sensations but rather as an essential element in the awareness of an experientially objective world. In the subsequent chapters, I argue that in being visually aware of the egocentric positions of the worldly objects around us we are often aware also of our own spatial locations with respect to them, and that the visual experience of the world around one and one’s own situation in it is often an essential component in the knowledge that a human agent will have of his own intentional actions.

Contents

<i>List of Figures</i>	ii
<i>Acknowledgments</i>	iii
<i>Introduction</i>	v
1. Can There Be Visual Perception Without the Visual Awareness of Space?	1
2. Vision, Self-Location, and the Phenomenology of the Point of View	17
3. Perception and Practical Knowledge	39
<i>Bibliography</i>	61

List of Figures

Figure 1: A representation of the Feature Integration Theory of visual attention	6
Figure 2: Visual illusion of self-motion (vection) induced by a rotating drum	26
Figure 3: The agent and observer models of action awareness	49
Figure 4: Non-observational, practical knowledge of one's intentional action	52
Figure 5: The Ebbinghaus illusion	59

Acknowledgments

The intellectual, personal, and professional debts incurred in writing this dissertation have been many, and most are no doubt evident from the text itself, but a summary accounting seems in order.

At various points during the actual writing process and the period of study that led up to it, Randall Amano, Karl Ameriks, Tony Bezsylko, Mike Caie, Bert Dreyfus, James Genone, Hannah Ginsborg, Alison Gopnik, John Keller, Sean Kelly, Niko Kolodny, Tania Lombrozo, Mike Martin, Steve Palmer, Matt Parrott, Bill Prinzmetal, Barry Stroud, Daniel Warren, and many others of my colleagues in my department's regular dissertation seminars were all especially encouraging or helpful in many different ways. Paul Franks, Bill Ramsey, and Fred Rush are three teachers from my earlier years of graduate school whose formative influence on my approach to philosophy, as well as to many of the ideas developed here, can hardly be overstated. A months-long e-mail exchange with Robert Briscoe was hugely helpful in getting me to see my way through some of the issues discussed in the final three chapters, and over the course of many seminars and countless hours of conversation James Stazicker has been a great friend and a challenging sparring partner whose model approach to philosophy has significantly influenced my own. And finally, I owe a tremendous debt to my wife Angela and to my thesis advisers, John Campbell, Alva Noë, and Lynn Robertson, who have been at once incisive critics and patient and encouraging mentors; it is hard to imagine how I could have completed this project without their support.

This dissertation is dedicated to my wife and children, and also to the memory of Susan Hurley, with whom I was able to discuss the project only in its most nascent stages but who would, I like to think, have found something in here to be characteristically excited about.

Space is not in the subject, nor is the world in space. Space is rather ‘in’ the world in so far as space has been disclosed by that being-in-the-world which is constitutive for dasein. Space is not to be found in the subject, nor does the subject observe the world ‘as if’ that world were in a space; but the ‘subject’ (dasein), if well understood ontologically, is spatial. And because dasein is spatial in the way we have described, space shows itself as *a priori*. This term does not mean anything like previously belonging to a subject which is proximally still worldless and which emits a space out of itself. Here “*apriority*” means the previousness with which space has been encountered (as a region) whenever the ready-to-hand is encountered environmentally. (Heidegger 1962: 146)

Introduction: Space and Self-Awareness

1 “Inner” and “Outer”

Hume scoured his mind for the source of his idea of himself, and his take on what he was supposed to be looking for quickly left him in quite a lot of philosophical trouble:

For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I can never catch *myself* at any time without a perception, and can never observe any thing but the perception. ... If any one upon serious and unprejudic'd reflexion, thinks he has a different notion of *himself*, I must confess I can reason no longer with him. All I can allow him is, that he may be in the right as well as I, and that we are essentially different in this particular. He may, perhaps, perceive something simple and continu'd, which he calls *himself*; tho' I am certain there is no such principle in me. (1978: 252)

What Hume sought was an impression of a thing possessed of “perfect identity and simplicity”, such that we “feel its existence and its continuance in existence” and is “not any one impression, but that to which our several impressions and ideas are suppos'd to have a reference”; for it is this, he supposed, that we have in mind when we think of “what we call our SELF” (1979: 251). And with its objective so defined it is hardly surprising that his search came away empty: even if he *had* turned up an impression whose intrinsic constancy corresponded appropriately to the purported simplicity of the self it is hard to see how it would have helped him much at all, as that would once again have been just another (albeit peculiar) perception, which could hardly have displayed in its appearance whatever relation it might have borne to the metaphysical subject supposed to be standing somewhere in the audience.

Hume did not take the metaphysicians' conception of the self very seriously, and in this much he arguably got things right; but the account he went on to give of the ordinary opinion of humankind was no less strange and metaphysical than the one he began by rejecting:

... setting aside some metaphysicians of this kind, I may venture to affirm of the rest of mankind, that they are nothing but a bundle or collection of different perceptions, which succeed each other with an inconceivable rapidity, and are in perpetual flux and motion. ... The mind is a kind of theatre, where several perceptions successively make their appearance; pass, re-pass, glide away, and mingle in an infinite variety of postures and situations. There is properly no *simplicity* in it at one time, nor *identity* in different; whatever natural propension we may have to imagine that simplicity and identity. (1978: 252-53)

Why, we might ask, if Hume really was out to give the *common* conception of what it is to be a person, did he not take his bearings from the observation that what we have in mind when we “suppose ourselves possess of an invariable and uninterrupted existence thro' the whole course of our lives” (1978: 253) is less the identity of our “selves” than of certain *human beings*? No doubt there are many available explanations, Hume's sly humor and the influence of Locke on subsequent British philosophical thought perhaps foremost among them; but certain features of

the surface grammar of English self-reference are also a likely culprit, as our ordinary talk of “myself”, “yourself”, “himself”, “themselves”, and so on can naturally suggest that we think of “selves” as peculiar kinds of entities that stand to be integrated into our commonsense ontology. Add to this the Cartesian background to modern philosophical thought about the mind and a generally dualistic account of personal immortality in which bodily continuity plays little if any role, and the question of what if anything could possibly justify us in thinking this way seems a pressing demand indeed.

But whether or not we reject some or all of these demands as illusory and insist instead that the thing I refer to when I speak or think of “myself” is not an immaterial substance but rather a bodily presence in the world, and so that I can “catch *myself* at any time” simply by looking down at my body or noting my reflection in a mirror, the philosophical problems of selfhood and personal identity are highly unlikely to go away; the roots of the metaphysical impulse run considerably deeper than that. Moreover, even when an appropriately naturalized picture of selfhood is set in place the traditional dualism of outward perception and inward retreat can still exert an influence on the accounts we go on to give of human self-awareness and its relationship to our awareness of the surrounding physical environment, as we find ourselves treating sensory channels dedicated to bodily interoception as the ecological counterpart to Kant’s “inner sense”, or “that by which the mind intuits itself or its inner state” (1933: A22/B37); they become the privileged source of knowledge of oneself, and we think of the “outer” senses as directed, if not exclusively, then at least primarily toward the things in the world around us. As José Luis Bermúdez describes the basic picture:

... the five senses are directed “outwards”—they are exteroceptive or exterosensitive, designed to inform us about objects and events in the world. They can, of course, be turned on oneself, as for example, when one looks at oneself through a mirror, but doing this provides a distinct sort of information about oneself, information that objectifies the body, failing to do justice to the sense in which the subject of perception is also the object of perception. This objectifying form of perceiving oneself is often contrasted with the form of self-perception from within, gained through what has been termed a ‘body sense’. Receptors in the skin, muscles, tendons, and joints, operating in conjunction with the vestibular system, yield proprioceptive information about bodily position that is crucial in orienting and acting within the world. This has led to a firm distinction in both operation and function, with the five exteroceptive senses deemed to provide information about the external world, while the proprioceptive system provides information about the self, in particular about bodily posture and movement. (1995b: 134)

Thought of in this way these privileged modes of self-awareness are not the ways a person intuits the “inner state” of his *mind*; rather, the thing whose states are made manifest through proprioceptive awareness is a living human being, and the bodily states of such a being are clearly spatial rather than, as Kant thought the contents of inner sense to be, merely a series of mental events arranged sequentially in time. But nevertheless such a conception leaves us room to distinguish the inner space of the perceiver’s body from the spaces occupied by objects in the surrounding world, and to hold that vision and other modes of ordinary sense experience are fundamentally our ways of finding out about spaces of the second sort, while the inner space of the body is revealed to us through different sensory channels. The primary purpose of this dissertation is to challenge such a division by arguing first for the importance of self-awareness

to the ordinary structure of visual experience, and then for the similar import of the contents of visual experience to the ordinary awareness of oneself.

2 Visual Self-Awareness

How should we think about the idea that visual experience might have as much a role to play in human self-awareness as in our awareness of the objects around one? A natural place to turn in motivating such a thesis is to the notion of “visual kinesthesia” developed by J.J. Gibson:

... vision is *kinesthetic* in that it registers movements of the body just as much as does the muscle-joint-skin system and the inner ear system. Vision picks up both movements of the whole body relative to the ground and movement of a member of the body relative to the whole. Visual kinesthesia goes along with muscular kinesthesia. The doctrine that vision is exteroceptive, that it obtains “external” information only, is simply false. Vision obtains information about *both* the environment *and* the self. In fact, all the senses do so when they are considered as perceptual systems ... (1986: 183)

On Gibson’s account of sense perception the chief experimental findings that help to secure this result have to do with the ways that organisms respond to certain patterns of optic invariance; the idea is that aspects of visual motion perspective alone can be shown, without contributions from non-optic sources, to have “told” a perceiver “not only about the earth but also about himself” (ibid.). As Shaun Gallagher puts the point:

Working together with the vestibular system, [visual proprioception and visual kinesthesia] help to distinguish between movements made by objects in the environment and one’s own movements. They can override vestibular information, however, and lead to a mistaken sense of movement. For example, when you are sitting on a train waiting for it to move and the train next to it begins to move, visual proprioception will provide a sense that you are moving. On the basis of such ecological information, outside conscious awareness, adjustments in posture are made in order to compensate for changes in the optical flow that accompany movement in the visual environment ... (2005: 45-46)

The Gibsonian account of visual kinesthesia has considerable psychological import, but given its focus on understanding vision as a “sensory system” rather than a source of conscious awareness it does not touch directly on an issue that is likely to be especially important to philosophers, namely that of the relationship between bodily self-awareness and the contents of ordinary visual *experience*. For one thing, as Gallagher’s account makes clear there is considerable room for optic information concerning the perceiver’s body to be processed subpersonally and to make little or no difference to the structure of conscious awareness: thus the behavioral and postural adjustments that organisms make on the basis of self-specifying visual information does not itself show that they can “*see* where they are going” (Gibson 1986: 183; emphasis added), where this is understood in terms of conscious visual experience rather than mere visual perceptual sensitivity.¹ Moreover, by focusing our attention just on those cases where self-motion is

¹ Thus Bermúdez writes that “Gibson’s position seems to be that conscious recognition is not implicated in ecological perception, although it might or might not develop out of ecological perception. It is perfectly possible for

perceived purely on the basis of how things strike the eyes the Gibsonian account somewhat obscures the fact that how things are presented in visual experience tends to outstrip the information available in the optic array: the fact that, as Gallagher suggests, vestibular information concerning the direction of gravity might have an influence on how things appear to a subject does not make that appearing a non-visual matter, any more than the illusory appearance of a flashing light produced by playing a discontinuous tone in a subject's ear (Shams, Kamitani, and Shimojo 2002) is not a genuinely visual illusion simply because of its non-optic determinants.²

By way of contrast, consider Alva Noë's account of the experience of being in an aircraft as its nose begins to rise:

At takeoff it will look to you as if the front of the plane, the nose, rises or lifts up in your field of vision. In fact, it does not. Because you move with the plane, the nose of the plane does not lift relative to you. No lifting, strictly speaking, is visible from where you sit. What explains the illusion of the apparent rising of the nose? When the plane rises, your vestibular system detects your movement relative to the direction of gravity. This causes it to look to you as if the nose is rising. The nose is rising, and it looks to you as if it is. But not for visual reasons. ... How things are experienced *visually* depends on more than merely optical processes. (2004: 26)

There are elements of this account that can be disputed; in particular, depending on how we think of the experience's spatial content it is not clear whether it will be right to think of this example as a visual "illusion" at all. What is crucial, however, is the distinction Noë draws between the notion of visual experience on the one hand (how things "look to you", how they "are experienced *visually*") and that of optic stimulation (or that which is "visible") on the other: there can be aspects of visual experience that outstrip how things strike the eyes, and this is made vivid in the fact that while the structure of the optic array remains unchanged as the plane takes off, the character of visual experience changes significantly. The question whether vision is a form of

a creature to have experience at the ecological level without any conscious recognitional capacities at all" (1995b: 163).

² There is confusion on this matter in Bermúdez 1998, where the issue of the possibility of visual kinesthesia is treated as equivalent to the question whether optic information alone can induce experiences of self-motion. Here, for example, is Bermúdez's discussion of the famous "moving room" experiments (Lishman and Lee 1973)

... subjects are placed on the solid floors of rooms whose walls and ceilings can be made to glide over a solid and immovable floor ... If experimental subjects are prevented from seeing their feet and the floor is hidden, then moving the walls backwards and forwards on the sagittal plane creates in the subjects the illusion that they are moving back and forth. This provides strong support for the thesis that the movement of the perceiver can be detected purely visually, since visual specification of movement seems to be all that is available. (1998: 111)

Bermúdez seems to be right in treating the moving room experiments as evidence that there is such a thing as the "purely visual" detection of the movement of the self, where by this is meant a detection of self-movement on the basis of information obtained by the eyes alone. He is wrong, however, in treating the possibility of visual kinesthesia in *this* sense as equivalent to, or even as entailing or entailed by, the claim that "the self has a place in the content of visual experience" (1998: 112): for the "purely visual" sense of the self brought out by the moving room experiments can (and indeed does) also give rise to *non*-visual experiences (like that of nausea or imbalance) in which the self seems to have a place; and furthermore there is no barrier to visual experiences with first-personal contents being affected by self-specifying information from non-visual channels.

self-perception in *this* sense is logically independent of the question that so exercised Gibson; it is a question about the metaphysics of human visual experience, rather than the nature of sensory processing in the human visual system. And it is primarily in Noë's sense, rather than Gibson's, that the question whether vision is a form of self-awareness rather than merely outer sense will be under investigation here.

3 Summary of the Argument

The argument of this dissertation is less a grand narrative than a series of philosophical-cum-psychological landscapes through which I hope to bring out a clearer view of the matters under consideration. We will be concerned first of all with how we should think about the role of spatial awareness in visual experience quite generally; then with the question of where self-awareness figures in an adequate account of the contents of visual awareness; and finally with the function that visual sense perception plays in helping to constitute the awareness of oneself. In what follows I will give a brief overview of the core argument of each of the chapters, together with a summary account of how those arguments should be thought of as interrelated parts of a thematically unified whole.

The topic of Chapter One is the deeply intuitive idea that the visual awareness of space is a necessary condition on the possibility of visual experience in general. This proposal, which I term the *Apriority Thesis*, finds its most famous articulation in Kant's claim that the intuition of space is that "in which alone the sensations [of outer intuition] can be posited and ordered in a certain form", a condition on representing sensations "as outside and alongside one another, and accordingly as not only different but as in different places" (1933: A20/B34, A23/B38), though I also discern compelling philosophical arguments for it in the writings of Edith Stein and the early Wittgenstein. On the interpretation I offer, the core claim of the Apriority Thesis has to do with the relationships between particular visible objects and properties and the visual experience of an experientially "absolute" space, which functions as a sort of experiential "container" that is a condition of the possible experience of the things that appear within it. In support of this picture, what proponents of the Apriority Thesis invite us to do is engage in imaginative exercises that bring out the limits of possible visualization and, through that, the essential structures of visual experience itself.

While the Apriority Thesis is clearly intended to express a metaphysical necessity that can be demonstrated on the basis of a priori philosophical reflection alone, I argue that certain such theses nevertheless generate empirical predictions that a posteriori experiment can put to the test: in this case, commitment to the Apriority Thesis requires us to predict that individuals in whom the visual awareness of space is extinguished or severely impaired should evince corresponding deficits in the visual perception of objects and spatial properties and relations. And notwithstanding the apparent impossibility of visualizing experiences that do not respect this constraint, I propose and discuss in some detail a neurological condition in which the constraint seems to be violated, in which the relevant aspects of visual spatial awareness remain intact even as the visual experience of an "absolute" space goes missing. Thus I conclude that this particular version of the Apriority Thesis deserves to be rejected, though I rebut the idea that this failing gives us reason to be skeptical of the worth of philosophical analyses of the nature of experience more generally.

In the concluding section of Chapter One, I briefly propose in place of the Apriority Thesis a different version of the claim that spatial awareness is required for the visual experience of spatial objects. What differentiates this proposal from the one rejected in the earlier portion of the chapter is that its emphasis is not on the possibility of “positing and ordering” visual *sensations*, but rather on the conditions for the possibility of what I call the *phenomenal objectivity* of visible shape: roughly, the idea is that part of what it takes for a material object in the field of view to have a visibly apparent structure corresponding to what P.F. Strawson calls the notion of an object “in the weighty sense” (1966: 173) is that it be experienced as having spatial properties other than its intrinsic shape, which can be experienced as changing while its figure remains visibly the same. This is a different idea from the more empiricist picture we find in the earlier version of the Apriority Thesis, and I argue that it provides a better fit with the empirical data and a more philosophically promising way to think about the way in which visual spatial awareness might be essential to visual experience in general.

Chapter Two continues the discussion of the metaphysics of visual spatial awareness by considering whether the spatially perspectival character of visual experience can be explained simply by taking its contents to be defined by an egocentric reference frame centered on the perceiver’s body, or whether instead we should think of visual experience as involving an awareness of the perceiver’s location as such. Having laid out the logical space that makes room for such a distinction I argue against a merely perceiver-relative conception of perspectival spatial content and in favor of what I call the *Self-Location Thesis*, according to which it is possible visually to experience one’s own spatial properties even when one’s body is entirely out of view. The argument of this chapter involves reconceiving Gibson’s notion of visual kinesthesia along the lines proposed in Section 2 above, as an aspect of the contents of visual consciousness rather than the structure of sensory processing; having suggested a way to think about whether self-locating contents are among the basic elements of how things visually appear to us I argue that views which try to avoid attributing such contents to visual experience fail to do justice to its qualitative character. I conclude the chapter by arguing more briefly for the further claim that self-locating contents are a *universal* aspect of human visual experience, and not just an element in how things visually appear in certain special cases. In both its modest and more ambitious forms, this argument for the Self-Location Thesis constitutes a key element in an account of visual experience as a form of self-awareness rather than purely outer sense.

In Chapter Three I make the case for a further thesis relating visual experience and human self-awareness, this time concerning the role of visual perception in intentional agency, and in particular in the essential ability of an agent to know what he is intentionally doing in a way that (i) cannot be shared by a second party and (ii) is not a way of knowing about his *unintentional* actions. Against views that attempt to explain these asymmetries by taking what Elizabeth Anscombe calls the “non-observational” character of our knowledge of our intentional actions to rest in an independence from sense perception I argue at length that we often do not, after all, have grounds for believing that we are actually doing what we intend that are independent of our experience of our actions as they unfold. It is true that an agent’s knowledge of his own intentional actions differs fundamentally from his knowledge of his unintentional actions and the intentional and unintentional actions of other people, but that distinctiveness needs to be accounted for in a different way than this.

What I go on to propose in the second half of Chapter Three is that we can best understand the distinctiveness of our knowledge of our intentional actions by recognizing that it is, as Anscombe puts it, a *practical* knowledge, where by this is meant a kind of knowledge that

is effectively involved in the actions that are its objects. When a person knows what he is intentionally doing his knowledge is distinctive not because it has a distinctive source, but rather that it plays a distinctive role in the agent's mental and behavioral economy: rather than being passive observers of our intentional actions, we instead bring our knowledge of those actions to bear in seeing to it that we bring about the things we intend. Thus our knowledge of our intentional actions is an element in the execution of those actions themselves; it is the cause of the very things it understands. In contrast to those views that take the self-knowledge that accompanies intentional action to be non-empirical, this account of the metaphysics of agency allows us to make sense of the special way in which we know what we are doing without beating the problematically inward retreat that loosens our handle on Anscombe's important recognition that human action is a worldly happening rather than a private event that takes place somewhere inside the agent's skin.

Thus I think of the three chapters of this dissertation as taking aim from two different directions at the philosophical tenability of the divide between inner and outer sense: first, by showing in Chapters One and Two that an account of self-awareness is essential to a theory of the visual experience of worldly space; and second, by showing in Chapters Two and Three the importance of visual experience to certain of the most fundamental aspects of the awareness of oneself. In place of the division this Introduction began with, we are thus invited to think of spatial awareness and self-awareness as fundamentally integrated aspects of our experience of how things are in the world.

Chapter One: Can There Be Visual Perception Without the Visual Awareness of Space?*

1 The Apriority Thesis

It is natural to think that the representation of space has a status in visual experience that is reasonably called “a priori” in at least the following sense: it is not possible visually to be aware of any object, property, or relation without experiencing it as positioned in visual space; and thus the visual awareness of space is a condition of the possibility of visual perception in general. Kant expresses such a view in a well-known passage from the *Transcendental Aesthetic*:

... in order that certain sensations be referred to something outside me (that is, to something in another region of space from that in which I find myself), and similarly in order that I may be able to represent them as outside and alongside one another, and accordingly as not only different but as in different places, the representation of space must be presupposed. The representation of space cannot, therefore, be empirically obtained from the relations of outer appearance. On the contrary, this outer experience is itself possible at all only through that representation. (Kant 1933: A23/B38)

The idea here is that perceptually representing the spatial distinctness of particular sensations requires the representation of a space that those sensations occupy distinct parts of; otherwise sensory items can be separated only in time, and there can be nothing deserving the title of “outer experience”. Similarly, Wittgenstein argues in the *Philosophical Remarks* that it is “obviously possible to establish the identity of a position in the visual field, since we would otherwise be unable to distinguish whether a patch always stays in the same place or whether it changes its place” (1975: 253-54). And the possibility of doing this is, he goes on to claim, not one that we find in visual experience only some of the time, but rather something that is required for there to be any visual experience at all:

If every point in visual space is marked out as distinct, then there is certainly a sense in speaking of *here* and *there* in visual space ... But is this property of having points marked out as distinct really essential to visual space; I mean, couldn't we imagine a visual space in which we could perceive only certain spatial relations but no absolute position? That is, could we picture an experience so? In something like the sense in which we can imagine the experiences of a one-eyed man? – I don't believe we could. For instance, one wouldn't be able to perceive the whole visual field turning, or rather this would be inconceivable. How would the hand of a clock look, say, when it moved around the edge of the dial? (I am imagining the sort of dial you find on many large clocks, that has only points on it, and not digits.) We would then be able to perceive the movement from one point to another – if it didn't just jump from one position to another – but once the hand had arrived at a point, we wouldn't be able to distinguish its position from the one it was in at the last point. I believe it speaks for itself that we can't visualise this. (1975: 254)

* I am especially grateful to John Campbell, Thane Naberhaus, Alva Noë, Bill Prinzmetal, Lynn Robertson, James Stazicker, and Daniel Warren for comments and discussions concerning materials related to this chapter.

Wittgenstein's immediate concerns in these passages have a somewhat different focus from Kant's, but there is a core idea that both philosophers seem to be committed to, namely that the possibility of visual experience presupposes the awareness of some sort of overarching space, as this latter sort of awareness is a condition on the possibility of "having points marked out" in visual perceptual consciousness. Finally, here is Edith Stein, in her elaboration of Husserl's notes on the "Systematic Constitution of Space":

... each body must have a location in intuition; each body is, as a matter of principle, given only as oriented. "In intuition," we said. Then we must immediately say that each body has its location, one that belongs to it. ... The possible system of locations or system of orientations, every location of each body, is thus comparable with every other one, and therefore each body has an orientation toward all the others, a location in relation to them. (Husserl 1997: 277)

Once again, we find here the idea that particular objects cannot be perceptually experienced except as positioned (located and oriented) in a single space and as standing in spatial relations to other (actual or possible) objects; there is no such thing as perceptual experience without an awareness of the twin "systems" of location and orientation. Clearly the ways that these philosophers articulate their commitment to what we can call the *Apriority Thesis* differ significantly from one another, not least in the fact that Kant's claim is meant to apply to spatial representation in general while Stein is talking specifically about perceptual intuition and Wittgenstein only about vision. But running through all three of these passages we can discern a shared commitment to the idea that *at least* in the visual domain, spatial awareness is a condition of the possibility of perceptual experience in general. My purpose in this chapter is to argue on empirical grounds that despite its undeniable philosophical attractiveness, the Apriority Thesis is simply false, and needs to be replaced by a different conception of the place of spatial awareness in visual perceptual experience. I will also make some general remarks about the epistemology of phenomenological analysis.

2 Two Varieties of Spatial Awareness

One philosopher who has more recently attempted a defense of something like the Apriority Thesis is Quassim Cassam, who argues in *The Possibility of Knowledge* for what he calls the "Spatial Perception Requirement":

(SPR) In order to perceive that something is the case and thereby to know that it is the case one must be capable of spatial perception. (Cassam 2007: 88)

According to Cassam the truth of SPR is a consequence of the fundamental status of spatial properties in the constitution of material objects; hence any perceptual episodes that can contribute to an understanding of the natural world "must somehow be connected to the ability to perceive [the] primary qualities" of things (2007: 120). If Cassam's argument succeeds, then it will have shown that a certain sort of spatial awareness is a condition of the possibility of sensory episodes that can contribute to empirical knowledge.

But the Apriority Thesis claims something more than the Spatial Perception Requirement. For according to Cassam the kind of spatial perception that is a precondition of empirical

knowledge “can either be understood as the capacity to perceive spatial properties or as the capacity to perceive space” (2007: 91). An Apriority Theorist, on the other hand, means us to focus precisely on the distinction that Cassam is here minimizing: Kant argues that we cannot represent things “as outside and alongside one another, and accordingly as not only different but as in different places”, without representing them as “in space”; Wittgenstein, that the visual perception of spatial relations requires the experience of “absolute position”; and Stein, that any physical object must be experienced as spatially positioned, and as standing in a range of spatial relations to other (actual or possible) objects. As Daniel Warren puts it, what Kant puts forward is the substantive hypothesis that in order to represent spatial properties and relations “a distinct representation, and thus, a distinct capacity is presupposed, namely a capacity to represent the spaces which objects occupy” (1998: 210). Thus the Apriority Theorist adds to SPR a further claim, namely that the possibility of spatial perception in Cassam’s sense, i.e. the perception of spatial properties and relations, requires the capacity for the representation of space. And this stronger requirement seems logically independent of Cassam’s.

Let’s call the kind of perceptual awareness corresponding to the sort of spatial representation that Cassam is concerned to reveal as a condition of the possibility of empirical knowledge “spatial awareness” (or “mere spatial awareness”), reserving the title “awareness of space” (or “awareness of space in the strict sense”) for the kind of perceptual awareness – i.e., the awareness of what Kant calls “the one all-embracing space” (1933: A25/B39) – that has just been argued to be the concern of the Apriority Theorists. If we read Kant’s talk of “the representation of space” as I have proposed, then clearly he is saying something more than Cassam’s argument will yield: on this understanding of the Apriority Thesis not only is spatial awareness required for empirical knowledge of the natural world, but mere visual spatial awareness itself requires the visual awareness of space.

3 Science and Metaphysics

How, then, is the Apriority Thesis to be argued for? According to Warren, the suppressed premise in Kant’s argument is an appeal to the nature of geometrical knowledge: it is because “[t]he ascription of spatial relations to objects presupposes many a priori claims about what combinations of spatial relations are or are not possible” (1998: 207), and because these claims are neither purely logical nor justifiable by mere experience, that the representation of space (or spaces) is required for that of spatial objects and relations. But such a premise assumes a discredited mathematical epistemology that will likely have little appeal to us, and makes it hard to explain what makes this way of thinking about spatial awareness so intuitively attractive. A twenty-first century Apriority Theorist will wish to proceed in a different way than this.

The passage quoted from Wittgenstein suggests a different route to the Apriority Thesis, one which appeals directly to the consequences of imaginative exercises that test the limits of possible visualization. For Wittgenstein, the question whether a given property is “essential to visual space” seems to be equivalent to the question whether we can “imagine” or “picture” a visual experience in which things are otherwise: the idea is that if it is impossible for us to *visualize* spatial relations without “absolute position”, or bodies that have no apparent location or orientation, then it follows that such experiences are, at least in our case, a simple impossibility. This sort of phenomenological method is clearly at work in the passage we cited from Stein as well, and it may be in the background of Kant’s thinking as well. Clearly this style of argument is dangerous when it is extended beyond the metaphysics of experience, and made to apply to the

structure of physical space itself: for without a lot of dubious metaphysical theorizing of the sort we find in Kant, we will have little good reason to claim that the necessary structures of the universe must reflect the way that we human subjects are (perhaps) necessarily constrained to perceive or otherwise represent it. But our way of interpreting the Apriority Thesis avoids this particular problem by limiting its application just to the structures of our particular kinds of minds, proposing that it is our capacity for visual imagination that gives us the kind of intimate relationship to the nature of visual experience that enables us to lay bare its essential structures from the philosophical armchair. By bringing to light facts about what we can and cannot visualize, the Apriority Theorist proposes to reach conclusions about what we can and cannot visually experience.

Understood in this way the arguments in favor of the Apriority Thesis are, of course, distinctively philosophical: they advance metaphysical claims about the necessary structure of conscious experience in ways that make no appeal to the outcome laboratory experiment or other a posteriori considerations. But this does not mean that Apriority Theorists are entitled to regard empirical data as altogether irrelevant to the critical evaluation of their view. For even if claims about metaphysical necessities cannot reasonably be supported by experimental means, they may nevertheless yield significant empirical predictions, and can therefore be put to the test by considering whether those predictions hold up. As José Luis Bermúdez puts it:

Neuropsychology can bring ... thought-experiments to life. It allows the formulation and testing of predictions. If a transcendental argument claims that the possibility of a subject's having a particular conceptual ability A is dependent upon his possessing a further conceptual ability B, then this clearly generates the prediction that A cannot exist in the absence of B. We can test such predictions by looking at what happens in neuropathies where ability B is severely impaired. If the argument is sound, we would expect the patient also to lack ability A. If, however, ability A remains unimpaired, then clearly the argument needs further examination. (1995a: 381-82)

In the case of the Apriority Thesis, the relevant prediction is that since the visual awareness of space in the strict sense is a condition of the possibility of visual experience in general, any individual whose capacity for the latter sort of awareness is extinguished or impaired should evince corresponding deficits in other aspects of visual experience, including the capacity for the “mere” spatial awareness of visible objects and properties. Since the Apriority Thesis asserts a metaphysical necessity, clearly its proponents cannot find real support for their view in particular cases where this sort of relationship is found to turn up; but by the same token, they will clearly be in trouble if its critics can find cases in which it does not.

4 Vision Without the Awareness of Space: A Case Study

With this in mind, consider Lynn Robertson's patient “RM”, who had suffered a pair of strokes that resulted in severe damage to his posterior parietal cortex, and whose capacity for visual spatial awareness was greatly impaired as a result. For example, when shown a display consisting of a screen with an ‘X’ at one of five locations along either the vertical or horizontal meridian and told to report whether the position of the stimulus on the screen was up, down, or center (in the vertical blocks) or right, left, or center (in the horizontal ones), RM averaged only 70% correct across all conditions. Similarly, when instructed to judge the relative position – left

or right in one block, up or down in the other – of an ‘X’ with respect to that of an ‘O’ that was also presented on the screen, RM was only correct approximately 50% of the time, a performance no better than chance (Friedman-Hill, Robertson, and Treisman 1995). RM could detect the stimuli well enough; he just couldn’t tell *where* they were, either on the screen itself or with respect to one another.

Moreover, RM’s deficits in visual localization were not a product of straightforward misperception, as if he routinely experienced stimuli as being to the top of the screen when they were in fact at the bottom, and so on. Rather, when instructed to report on the location of a visually presented object RM “had to be prodded to guess”, and “would shake his head back and forth and protest that he did not know where the word was” (Robertson 2004: 170). As Robertson puts it:

During early testing of his extrapersonal spatial abilities he often made statements like, ‘See, that’s my problem. I can’t see where it is.’ He also found it hard to describe what his perception was like. His explanations suggested that objects that popped into his view were not mislocated per se. Rather, they simply had no location in his perceptual experience. (Robertson 2004: 158-59)

So on RM’s own account, he (1) could see only one item at a time, and (2) experienced those things as having no location or orientation in visual space. This seems as good a test case as any for the sorts of conclusions advanced by the Apriority Theorists so far canvassed: for it was precisely the capacity visually to experience particular things as positioned either in visual space itself or with respect to other visual particulars that had been extinguished in his case. Thus the crucial question for our purposes is whether RM’s remaining visual perceptual capacities were impaired in the sorts of ways the Apriority Thesis predicts.

4.1 A Causal Role for Spatial Awareness?

A recent paper by John Campbell interprets the case of RM in a way that would seem to confirm at least some of the Apriority Theorist’s predictions, arguing that RM had “impaired object perception” (2007: 550), and that we can attribute this impairment to his visuospatial deficits. The key piece of evidence Campbell cites for his view is the way in which RM’s ability to detect target stimuli was undercut when their status as targets was defined by the possession of *conjunctions* of properties along different featural dimensions. For instance, when presented with displays consisting of two colored letters and asked to give the name and color of the first letter he saw, RM reported “illusory conjunctions” – in which the color of one object was experienced as conjoined with the shape of the other – at least 13 percent of the time, even with display times of as long as ten seconds (Friedman-Hill, Robertson, and Treisman 1995). He perceived the features (in this case, the colors and the shapes) that were on the screen just fine, but misperceived the objects those features were properties of. And since, according to Campbell, to perceive an object is to perceive “a single thing with many properties” (2007: 550), he concludes that RM’s impaired capacity for visual spatial awareness had causal ramifications on his visual perception of individual objects.

Following Robertson, Campbell proposes that we can get an initial handle on why there might be such a causal relationship by drawing on the notion of a visual “master map”, as it figures in Anne Treisman’s Feature Integration Theory of visual attention (Treisman and Gelade

1980). According to this well-known approach to the visual binding problem, different sorts of visible features – colors, shapes, orientations, and so on – are detected by the human visual system in an early and automatic fashion and represented in separate visual “feature maps”. Spatial representation in the feature maps is only very coarse-grained, reflecting the isomorphic spatial topography between information projected from the eye to the neocortex. The need for visual space arises, however, when two features must be combined into a single percept in order to, for instance, perceive an ‘X’ as red and an ‘O’ as blue. Visual attention, by “spotlighting” a particular location in visual space, leads the visual system to check all the various feature maps that have been stimulated by the visual array and figure out which properties are at that location; thereby recombining them into a unitary representation of an object.

Considerable evidence has been provided to support such a view, one instance of which will be sufficient for our purposes. Imagine you are asked to report the presence of a single target stimulus from within an array of non-target elements: Feature Integration Theory predicts that if target status is defined by the possession of a single property (such as color or shape alone), then automatic “feature-detection” will suffice, and you will be able to detect the target immediately and so without interference from the total number of non-target elements. If, on the other hand, the target is defined by the possession of a *conjunction* of two or more distinct kinds of properties (being a red ‘X’, say), then focal attention to its location will be required to detect it, and so response times will increase linearly with the number of objects available to be scanned. And in normal subjects, this is exactly what we find.

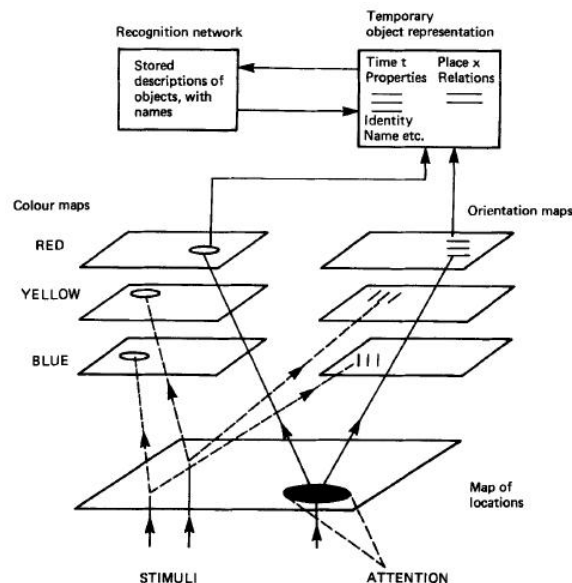


Figure 1: A schematic representation of Feature Integration Theory (Treisman 1988: 202).

Crucially, however, Campbell builds on this theory by proposing that it is specifically the *conscious* visual awareness of location that is required for feature-binding to occur. For example, when RM was shown a display consisting of a vertically oriented rectangle with one of the words ‘UP’ or ‘DOWN’ presented at its top or bottom and asked to (1) read the word and (2) report which end of the rectangle it was located in, his responses exhibited the Stroop effect (Stroop

1935): despite being unable to say where the words were located, RM was 142 ms slower to report the identity of the word in the “incompatible” conditions (where ‘UP’ was at the bottom of the rectangle, or ‘DOWN’ at the top) than the “compatible” ones (Robertson, Treisman, Friedman-Hill, and Grabowecky 1997). This shows that the locations of the words with respect to the rectangles were represented somewhere in RM’s visual system, and that this information had a causal impact on his verbal reports. Nevertheless, he clearly was not *aware* – consciously, that is – of the words’ locations: as noted above, when asked to report where they were RM insisted that he did not know, and “had to be prodded to guess the location” (Robertson 2004: 170); and when he did come to hazard a guess, it was accurate only 51% of the time (at chance, for a two-alternative forced-choice task). RM’s intact “implicit” visual representation of spatial information – as opposed to the “explicit” grasp of it that would have come with the conscious visual awareness of spatial location – was not enough on its own to enable the veridical binding of perceived properties into unitary representations of individual objects.¹ Thus, Campbell concludes, deficits in visual spatial awareness make for deficits in the visual perception of objects, much as the Apriority Thesis predicts.

But in fact things are not that straightforward. For despite his inability to experience how visual stimuli were located and oriented in space, *another* type of visual spatial awareness seems to have been perfectly intact in RM, as illustrated for example by the results of a letter identification task: he was shown a single letter, and asked to report what he saw. In a block in which the stimuli consisted of the letters ‘b’, ‘d’, ‘p’, and ‘q’, RM correctly identified the letter on 23 of 32 trials, and five of his nine errors consisted in confusing mirror-image pairs. Similarly, in a block whose stimuli were the letters ‘o’, ‘e’, ‘c’, and ‘d’, RM made only one error in 16 trials, and that was to read the ‘d’ as a ‘b’, another case of mere mirror inversion (Robertson & Treisman 2006). What is remarkable about this finding is that, just as in the studies described earlier that tested the awareness of location or the binding of features into wholes, RM did have veridical experiences of the *intrinsic shapes* of the display items, and so of how their proper parts fitted together to form spatial wholes. He did not, for example, experience the straight line in a ‘d’ as extending one side of the curve, as in a capital ‘J’, or the sideways line in an ‘e’ as running vertically through the curve, as in a ‘ø’ sign; rather, his visual errors were limited to illicit rotations or mirror-image reflections. In this limited sense, then, his visual spatial awareness was perfectly intact. We might say that RM was veridically aware of the bit of space that the figure took up (an “object-space”, we can call it), but regularly misperceived – or perhaps did not have any conscious experience at all of – how that shape was oriented in the larger space in which it was situated.

So it is important not to overstate the extent of RM’s visuospatial impairments (as Campbell seems to, when he simply says for instance that RM “had no spatial awareness” (Campbell 2007: 550); compare this to Robertson’s somewhat more cautious description of his condition as one in which there was “no ‘there’ there” (Robertson 2004: 6)). As Robertson and her colleagues put it, it was RM’s visual experience of *extrinsic* or *between-object* spatial properties that had gone awry, while that of *intrinsic* or *within-object* spaces and spatial relations was pretty much intact well in his case. This is why he could perceive the shapes of the letters up to isomorphism, but frequently misperceived how those shapes were oriented in space. The

¹ To say this is not, however, to foreclose the possibility that more careful empirical investigation might reveal the possibility of a “blindsighted” solution to the binding problem, by turning up subjects in whom the awareness of location is impaired but who can nevertheless detect conjunctions of features with a high rate of success. (Here I am grateful to Erica Klempner.)

visual awareness of space has two distinct “dimensions”, as it were, only one of which is relevant to the way that features such as shape and color are bound together into the visual perception of a unified object.

How does this finding square with Campbell’s claim that conscious spatial awareness is causally implicated in visual object perception? Consider the following quote from Robertson and her colleagues, in which the capacity to conjoin different kinds of features into a veridical representation of an object seems to be regarded as far less relevant to what they call “object perception” than the awareness of intrinsic shape itself:

Patients with Balint’s syndrome can identify a perceived object. *That is*, they are able to perceive a defining shape. Because of this ability, it was assumed that they saw the object with all its features intact, but this turns out not to be the case. The assumption also led to the proposition that spatial attention was reduced in those patients to cover only the spatial extent of the object perceived. However, this does not capture the full dimensions of the problem either. We have shown that such patients have abnormal feature binding as well. They miscombine a feature such as colour, size, or motion in the scene with the shape they perceive. (Friedman-Hill, Robertson, Desimone, & Ungerleider 2003: 424; emphasis added)

In this quotation object perception seems to be regarded as one thing, while feature binding is quite another. The former is thought to rest crucially on the veridical experience of the shape of a thing, while the latter is just a matter of attaching other properties to a thing that has – in virtue of the successful recognition of its shape – already been perceived. And so on Robertson’s way of speaking, the visual awareness of space in the strict sense appears causally irrelevant in the visual perception of objects, since as we have already seen the perception of a thing’s “defining shape” is altogether independent of the awareness of its location and orientation.

This latter way of thinking about what is required for object perception is the one which, by and large, Robertson and her colleagues rely on in describing RM’s condition. They say, for example, that he “is frequently aware of only one object in the visual field or one group of clustered objects” (Friedman-Hill, Robertson, Desimone, & Ungerleider 2003: 4264); similarly, they write in another paper that “When shown two objects he often only saw one. When he did report both, he did so slowly, and seemed to see them sequentially” (Robertson, Treisman, Friedman-Hill, & Grabowecky 1997: 297). And these quotations are by no means uncharacteristic; in general, RM and other Bálint’s patients are consistently described in the neurological literature as individuals who perceive objects but are unaware of their locations and sometimes misperceive certain of their features.

We can find some philosophical support for this way of thinking in the idea that the defining shape of a physical object occupies a special status among the properties that make it the kind of thing it is: in the same way that, as Cassam puts it, “the capacity to perceive their primary qualities is ... a necessary condition for the perception of material objects because it is in virtue of their primary qualities that they count as material objects in the first place” (2007: 124), so the perception of a thing’s primary qualities may also be sufficient to perceive that thing. This line of argument is further strengthened when we consider that not only is a thing’s defining shape a primary quality, but it is also a property *intrinsic* to the thing that has it: that is, it is an aspect of the thing that is independent of its relations to anything else, including (if indeed it makes sense to think this way) the larger space that surrounds it. Thinking about things in this way, and compiling for argument’s sake the union of all of Locke’s various lists of

primary qualities (see Alexander 1985: 133-34), we can that among *solidity, bulk, figure, mobility, situation, number, texture, and motion of parts*, it is because figure is both (i) a primary quality that (ii) is intrinsic to the thing whose figure it is, that the perception of a thing's defining shape suffices for the perception of that object.

If this is the criterion for visual object perception that we adopt, then Campbell's conclusion rests on a misdiagnosis of what has gone wrong in RM: it is not that his capacity visually to perceive *objects* has been impaired by his impaired visual awareness of space, but rather his capacity for the veridical visual perception of which properties those objects have. But as we have seen, according to Campbell the reason his stronger conclusion is warranted is that object perception requires more than the perception of mere features; rather, it requires the perception of how those features are combined into single objects. As he puts it:

His [RM's] object perception was impaired, *in the following sense*. He could not identify which conjunctions of features were present in an array: for example, he could not tell whether a red square was present in an array of variously colored circles and squares. Nonetheless, he was relatively good, well above chance, at saying which features were present in an array, even though he could not reliably say which features were conjoined with which. So he could tell whether redness was present and whether squareness was present, even though he could not reliably tell whether a red square was present. If we think of an object as a single thing with many properties, we can describe the situation by saying that RM had impaired object perception but intact feature perception. (2007: 550; emphasis added)

Campbell can describe things in this way, however, only by insisting that the visual perception of intrinsic shape does not already qualify as the perception of "a single thing with many properties", as opposed to a mere bunch of features present in an array. And there is good reason to think that any such insistence will be mistaken: for example, in the letter identification task described earlier where RM correctly identified a lowercase 'e', it would have been impossible for him to do this had he not perceived the very same object both *as having a curved side* and also *as having a horizontal line through the middle*. Even if there were others of the figure's properties – its size, its location, its orientation, and perhaps even its color – that he failed to perceive, the content of his visual experience still involved something analogous to the ascription of multiple predicates to a singular referent: he perceived a particular thing as having a complex shape that consisted of multiple figural aspects. And so it seems that we should say that even Campbell's own requirements for visual object perception were satisfied in this case, and that the mere perception of a structurally complex shape suffices for the perception of an object.

Importantly, this sort of featural complexity is a quite general characteristic of the way shapes are presented in visual experience, and the idea that a kind of binding together of visible properties is required for the perception of intrinsic shape is an important part of the psychological literature. For example, Treisman 1996 distinguishes seven different types of binding relevant to visual object perception, among them *part binding*, in which "the parts of the object [are] segregated from the background, and bound together" (1996: 171). That RM consistently managed veridically to perceive the intrinsic shapes of letters and even words clearly suggests that there is something which, even on Campbell's own terms, deserves to be called a kind of object perception, and which was perfectly intact despite his visuospatial deficits. In any case, it seems wrong to say that all RM was able to do visually was to "say which features were present in an array": he could do this, but of some of these features he could *also*

tell “which features were conjoined with which”, and moreover how those features fit together into unitary spatial wholes. Even on Campbell’s own terms, we do not have sufficient grounds to conclude that RM’s visual impairments impacted his capacity to meet the minimum conditions necessary for visual object perception. The condition of RM ultimately contradicts, rather than supporting, the claim that the visual awareness of space is even so much as causally relevant in the visual perception of objects.

4.2 Spaces and Shapes: RM and the Apriority Theorist

Insofar as Campbell’s primary objective is to argue for the claim that there is *some* causal role in visual perception that is played by conscious spatial awareness, the foregoing criticisms of his argument are ones that he can largely take on board: for even if RM’s impairments in the visual awareness of space did not undermine his capacity for the visual perception of the intrinsic shapes of things, Robertson’s description of his condition does indicate that his visuospatial impairments led to deficiencies in the way his visual system bound together distinct sorts of visible features. When we turn to the Apriority Theorist, however, the finding that RM had intact capacities for visual shape perception raises some much more fundamental questions. For what the Apriority Thesis proposes is a sort of *experiential Newtonianism*, according to which the visual awareness of space serves, like physical space on Newton’s view, as a sort of experiential “container” within which spatial particulars need to be placed if they are going to be made visually apparent to us at all. According to the Apriority Theorist, there can be no such thing as visible spatial structure unless visual *space* is experienced as well: for what is there to the spatiality of an object aside from the room it takes up in space itself, and how could something be visibly spatial without appearing somewhere in visual space? But in RM’s condition, it was precisely an aspect of “mere” spatial awareness, i.e. the visual awareness of the intrinsic spatial qualities of physical objects, that remained intact despite his incapacity for the visual awareness of space in the strict sense, i.e. of where those objects were positioned in any sort of larger spatial framework. Whereas the Apriority Thesis treats the awareness of within-object spatial relations as a special case of the awareness of the spatial relations that hold between distinct things, and so the visual experience of shapes and particular spaces as derivative upon the experience of space itself, RM’s condition indicates that these perceptual achievements are entirely distinct.

Is there a way to interpret RM’s condition that avoids this startling implication? The natural response for the Apriority Theorist is to treat RM’s visual awareness of intrinsic shape as *itself* an instance of the visual awareness of space in the strict sense: the thought would be that thanks to the fact that he could experience only one object at a time RM was subjected to a sort of “tunnel vision”, but that in visually perceiving what we earlier called “object-spaces” he thereby counted as visually aware of space in the sense that the Apriority Thesis requires. For according to the Apriority Theorist, simply being able to relate visual sensations so as to distinguish, say, an ‘e’ from a ‘ç’ sign demands the visual awareness of the spaces that those objects occupy; thus RM’s condition is compatible with the Apriority Thesis after all.

For an orthodox Kantian, the claim that RM was visually aware of space in the strict sense will be thought to find further support in the fact that, as noted above, RM could differentiate members of mirror-image pairs at a rate better than chance: he correctly identified the stimulus letter from the range ‘b’, ‘d’, ‘p’, and ‘q’ almost 75% of the time, and made only four errors that did not involve mere reflections or rotations. According to Kant, the possibility

of differentiating “incongruent counterparts” in this way requires the awareness of how they are related to space “as a whole”:

What indeed can be more similar to, and in all parts more equal to, my hand or my ear than its image in the mirror? And yet I cannot put such a hand as is seen in the mirror in the place of its original; for if the one was a right hand, then the other in the mirror is a left, and the image of the right ear is a left one, which can never take the place of the former. Now there are no inner differences here that any understanding could merely think; and yet the differences are inner as far as the senses teach, for the left hand cannot, after all, be enclosed within the same boundaries as the right (they cannot be made congruent), despite all reciprocal equality and similarity; one hand’s glove cannot be used on the other. What then is the solution? These objects are surely not representations of things as they are in themselves, and as the pure understanding would cognize them, rather, they are sensory intuitions, i.e., appearances, whose possibility rests on the relation of certain things, unknown in themselves, to something else, namely our sensibility. Now, space is the form of outer intuition of this sensibility, and the inner determination of any space is possible only through the determination of the outer relation to the whole space of which the space is a part (the relation to outer sense); that is, the part is possible only through the whole, which never occurs with things in themselves as objects of the understanding alone, but does occur with mere appearances. We can therefore make the difference between similar and equal but nonetheless incongruent things (e.g., oppositely spiraled snails) intelligible through no concept alone, but only through the relation to right-hand and left-hand, which refers immediately to intuition. (1997: 38)

Leaving aside the idealist baggage, the central claim Kant is arguing here is that the “inner difference” between, say, a right hand and a left hand or a lowercase ‘p’ and ‘q’ cannot be accounted for solely in terms of the spatial relations between the parts of things; we need instead to consider how those things are oriented within space itself. On this picture, the fact that RM could recognize intrinsic shapes and that mirror-image pairs were generally distinguished in his visual experience shows that he had a residual awareness of space in the strict sense after all: he experienced things as having spatial orientations, though not as spatially located.

But this way of thinking about RM’s condition requires us to conflate two different aspects of visual spatial representation, namely the experience of one part or another of an object as the top (or bottom, or right or left) *of that thing*, and the experience of such a part as oriented upward (or downward, or rightward or leftward) *in space*. It is true that in perceiving a stimulus as, say, a ‘p’ rather than a ‘q’ there is a sense in which RM would experience the curve as falling on the right side of the line rather than the left, but this is not yet to say that he experienced the *letter* as having any orientation in space. And crucially, the fact that RM “always saw things as upright” (Lynn Robertson, personal communication) gives us good reason to resist saying exactly that: for surely part of what it is to experience something as upright rather than upside-down is for there to be a visual possibility of experiencing it otherwise than that; RM’s visual awareness of the tops, bottoms, rights, and lefts of particular *objects* should be held carefully apart from the idea that he experienced those things as situated in “an oriented *space*, a *space* in which there is an above and below and a right and a left” (Wittgenstein 1975: 255; emphasis added). Whereas the Apriority Thesis treats the awareness of particular shapes as a special case of the awareness of space itself, RM’s condition indicates that at least in the visual domain these

perceptual achievements are quite distinct, no matter how unimaginable this may be from our perspective.

5 Some Methodological Observations

What does this discussion suggest in general about the philosophical promise of armchair reflection of the sort that we took to motivate the Apriority Thesis? Cassam glosses what he calls explanatory *minimalism* as the denial “that there are substantive enabling conditions for the acquisition of perceptual knowledge that, unlike physiological and biological conditions, can be established without any empirical investigation” (2007: 87). Clearly the Apriority Thesis proposes a condition on the possibility of perceptual experience of exactly the sort that a minimalist will reject out of hand: so does the fact that we seem to have proven it wrong by appealing to empirical considerations show, or provide evidence for, a minimalist position more generally? Must we abandon hope of the possibility of non-empirical knowledge of the necessary structures of consciousness and human knowledge, instead ceding these subjects to a purely experimental inquiry?

Across the philosophical spectrum from the minimalist’s radical empiricism we can juxtapose an staunch *conservatism* about the conclusions of a priori investigation, according to which they are to be upheld in the face of a posteriori scrutiny almost come what may. On this view the positions we arrive at by way of phenomenological inquiry and other modes of armchair analysis do indeed generate empirical predictions, but it is the responsibility of the philosopher-scientist to reinterpret initially recalcitrant empirical data to bring them into line with common sense. Thus a conservative who is deeply persuaded by the kinds of phenomenological analyses described in Section 3 might argue that our inability to imagine visual experiences anything like those that RM claims to have had – e.g., ones in which objects “simply had no location in his perceptual experience” – gives us license to be skeptical of RM’s ability to describe what his visual experience was like: perhaps he couldn’t *report* on where things visually seemed to him to be, but that doesn’t mean that there wasn’t a fact of the matter about how they were positioned in visual space; or perhaps we should push back against the idea that he really enjoyed conscious visual experiences of the intrinsic shapes of things.² We should not allow the subjective reports of someone like RM to undermine our basic philosophical convictions about the scope of visual possibility.

There are many things that can be said in response to this sort of position, but two points seem especially apposite. The first is that, conceding for a moment the intuitive ground on which the conservative’s objection is based, it is not as if the conceptual distinction between the awareness of space and the awareness of particular shapes has no echo at all in our visual phenomenology: for example, you can experience the very same object first as upright and then as lying on its side, but without experiencing any change at all in the intrinsic shape it looks to have. This clearly indicates that the visual awareness of within- and between-object spatial properties cannot simply be reduced to visual spatial awareness of a single basic sort, as an especially clumsy version of the Apriority Thesis might have it; and so while this observation certainly does not entail that these two aspects of visual experience can operate entirely apart from one another, it at least provides an ordinary experiential analogue of the independence claim that the case of RM invited us to draw.

² I am grateful to Daniel Warren and Alva Noë for pushing me to take this sort of position more seriously.

Secondly and more importantly, there is good reason to resist the idea that the limits of the imagination are a perfect guide even to the limits of possible experience: a sufficiently ignorant and unimaginative person raised in a thoroughly grayscale environment, for example, may find it inconceivable – may find himself unable to “picture” the possibility – that one might enjoy a huge range of visual experiences of a sort hitherto unknown to him. Thus in the same way that, say, mathematicians’ postulations of non-Euclidean spaces forces a revision in our intuitive understanding of the conditions under which lines can and cannot intersect, so the conditions of individuals like RM challenge us to reconsider even our most basic phenomenological intuitions, by presenting us with individuals who, though like ourselves in a lot of important ways, have minds that differ along significant dimensions from our own. Of course we should be cautious in the way we interpret their subjective reports and draw out their philosophical consequences, but in general it seems wise to appeal to further empirical investigation, rather than more first-personal phenomenology coming from our own side, if we want to show that what such individuals say about their own experiences simply cannot be right. And with RM, whose verbal spatial vocabulary and capacities for spatial reasoning and spatial perception in non-visual modalities were perfectly intact, there is little reason to doubt what he says about the nature of his visual experience: his claims about what visual experience is like in his own case seem to be on stronger ground than our intuitions about what is and is not visually possible.

This is not, however, to concede the dialectical ground to the minimalist, and so to give up the possibility of acquiring substantive knowledge of the essential structures of the mind through non-empirical means. For phenomenological exercises and other sorts of a priori investigation can clearly provide a *kind* of evidence for claims about, say, what visual experience must be like even if that evidence is not, as conservatives like Kant and the classical phenomenologists would have had it, perfectly apodictic and so immune to empirical scrutiny. What the case of RM teaches us is not to give up altogether on the possibility of substantive non-empirical knowledge, but rather to treat the fruits of such inquiries like any other instances of putative a priori knowledge: they may warrant a place somewhere near the center of our web of belief, but are not thereby immune to revision in the light of what we learn down the line. It is by being bold in articulating the empirical phenomena that their a priori conclusions require them to predict, and meticulous in seeking out and scrutinizing findings that can disconfirm them, that practitioners of non-empirical methodologies can make genuine contributions to the study of the mind.

6 Phenomenal Objectivity

According to the argument of Section 4, RM’s case provides evidence for a strong dissociation between the visual awareness of space and the visual awareness of particular spatial objects, in the following sense: it is possible visually to be aware of a complex shape without being aware of the space in which it is situated; and in being aware of a complex shape you are aware of the defining properties of a physical object. But there are ways of thinking about what object perception requires that make the mere awareness of a bundle of defining properties seem to fall well short of the bar; thus Kant writes in the *Critique*:

Everything, every representation even, in so far as we are conscious of it, may be entitled object. But it is a question for deeper enquiry what the word “object” ought to signify in

respect of appearances when these are viewed not in so far as they are (as representations) objects, but only in so far as they stand for an object. (1933: A189-190/B235-236)

Kant's idea here is that some of the things that are experientially presented before the mind have a sort of dual nature: in addition to being the ways they are, they also manifest the states of objects other than themselves. Hence perceiving, in this sense, is not just being saddled with a mass of mere sensory affections ("a blind play of representations, less even than a dream" (Kant 1933: A112)), but instead puts the perceiver in touch with how things are *in the world*, not just in some intra-mental domain whose layout is constituted entirely by the nature of his mental states. As Strawson writes, understood in this way the notion of an object "is to be taken more weightily than we might at first have thought", as it "carries connotations of 'objectivity'" (1966: 73). In this concluding section I want to sketch out a way in which this weightier conception of what it is to be an object can reveal a way in which the awareness of space is crucial in visual object perception after all.

In the hands of Strawson and his followers, Kant's "weightier" conception of objecthood has largely been put to work in exploring the conditions of the possibility of objective thought. Thus Strawson starts off his discussion of Kant's view of formal logic with a reminder that "all experience, all empirical knowledge, requires the co-operation of sensibility and the understanding, i.e. the awareness of particular objects as falling under general concepts"; and therefore, he continues, experience of objects in the weighty sense "is identical with making *judgements* about objects" (1966: 74). But if we think of perceptual experience as something that is fundamentally belief-independent, as a way of being in touch with the world that is ultimately *prior* to our conceptualization of it in explicit judgments, then the idea of an object in the weighty sense may still have a role to play in our phenomenological theorizing: the crucial task here will be to offer an account of *what it is like* to experience things that are more than mere mental representations, things that are present before the mind but nevertheless appear to have an existence outside the mind.

We can call the aspects of experience in virtue of which it presents us with objects in this weightier sense the characteristic of *phenomenal objectivity*. The crucial feature of phenomenal objectivity is that there is room for a distinction within the structure of experience itself between the appearances of things and how those things appear to be: if something is phenomenally objective, then its *esse* is not *percipi*; it will be possible for there to be changes in the perception of it that are not changes in the way it is perceived to be. Put somewhat differently, the idea behind the notion of phenomenal objectivity is that in order to experience something as an object in the weighty sense it is not enough simply for there to be some "object = *x*" that is "distinct from all our representations" and therefore "nothing to us" from the perspective of sensory consciousness (Kant 1933: A104), but which our perceptual episodes somehow manage to "stand for"; rather, we want to understand how a perceptual experience's being an episode of awareness of some mind-independent object is reflected in the subjective character of experience itself, and is more than a brute relation between perceptual experience and something outside the mind.

In the case of the visual awareness of shape, the crucial factors that constitute the possibility of phenomenal objectivity have to do with what is sometimes called the "two-dimensional" structure of visual shape constancy. The first dimension of shape constancy consists in the fact that when you visually perceive a spatial object – a penny, say – and the spatial relations between you and it change in such a way that there are consequent changes in way it stimulates your visual system, in ordinary circumstances you will nevertheless experience

the shape of the object as *unchanging*: if for example the penny was initially upright with respect to you and now is tilted backward, the fact that the pattern of retinal stimulation it gives rise to will first be (very roughly) circular and then (very roughly) elliptical will not cause you to experience the penny first as having one shape and then as having another. Nevertheless, in a case like this there is going to be *some* sort of change in the way the penny's shape appears to you: thus the second dimension of visual shape constancy consists in the fact that when a thing changes its position with respect to your eyes and you experience its shape as constant there is an important way in which your experience is changing. And it is due to this possibility that the visual experience of shape counts as phenomenally objective: there is more to an apparent shape than the particular appearances of it, as the way a shape appears to be is not reducible to appearance it has at any given moment.

What has to hold true of visual experience in order for shapes to be experienced as phenomenally objective? The standard account of visual shape constancy derives from the work of Irvin Rock, who in describing what he calls the “all-pervasive tendency” of perception “toward constancy of object-properties despite variation of or difference between the proximal stimuli” (1983: 24) acknowledges that perceptual constancy has the sort of the two-dimensional structure described just above. Here is how Rock suggests that we can account for the experiential complexity of visual shape constancy:

A circle seen from the side, let us say at a 45-degree angle, may in one respect be said to continue to look circular, ... but its elliptical retinal image is not without some perceptual representation. Again instructions to match in accord with the “projected shape” rather than the objective shape will result in matches that are closer to the shape of the retinal image than to the object shape. Although it is difficult to describe the nature of this aspect of shape perception, perhaps the term *extensity relations* will suffice. We are aware that one diameter of the circle has a greater extensity in our field of view than the other while nevertheless simultaneously experiencing the objective sizes of these diameters as equal. (1983: 256-57)

On this view, what we have been calling the “appearances” of things are identified with a residual awareness, or “proximal-mode experience” (1983: 254), of patterns of stimulation at the sensory periphery: thus the “elliptical retinal image” cast by the penny is what explains its appearing differently when tilted with respect to the observer than when it is upright. The ways that things *appear to be*, on the other hand, are the product of the familiar sorts of computational processes postulated by the cognitive science of vision, which factor out the distortive aspects of incoming stimulation to arrive at a veridical representation of the distal world. And since such stimulus patterns can change (or remain the same) even though the corresponding worldly objects remain constant, this distinction between proximal and distal (or “world”) aspects of awareness – both of which, Rock argues, must coexist in any case of perceptual experience even though we may be more attentive to one aspect (usually the distal one) than the other – is supposed to explain how outer objects can appear to remain the same despite changes in the ways they appear.

But it is hard to see how this proposal can do justice to the visual experience of shape constancy. On pain of violating Leibniz's Law, the objects of awareness of the proximal- and distal-mode experiences – the elliptical shape in the visual field, on the one hand, and the circular shape in the external world, on the other – must be distinct from one another; but how can it be that a change in the appearance of one shape – in this case, a distal circle – can be accounted for

simply by changes in the experience of an entirely *different* thing – here, a proximal pattern of extensity relations? Granting for the sake of argument the distinction between proximal and distal modes of awareness, there is no question that as the circular object tilts or otherwise changes its viewer-relative position the image that it casts on the retina, and so also the corresponding pattern of extensity relations in the visual field, will change accordingly; and as such, there will be a change *in the domain of* perceptual “appearance”. But not all such changes will be changes in the appearance *of the circle* that appears (“distally”) to be constant: guaranteeing this kind of coincidence requires something more than the resources Rock’s view puts at our disposal.

Thus what we are after is an understanding of how the two “dimensions” of perceptual experience that are required for phenomenal objectivity can be experientially unified so as to count as different aspects of the experience of one and the same thing: it is the shape of the *very same* penny, for example, that both changes in appearance and appears to remain the same. And surely the most natural way to explain this is by noting that when an object changes its viewer-relative position but is experienced as having the same shape, there are aspects *other than* its intrinsic shape that are experienced as changing: and crucially, the aspect of a thing’s intrinsic shape that seems to affect its visual appearance most directly is its spatial position with respect to the viewer. When the penny moves from being upright to being tilted, we can’t say all that there is to be said about the visual experience of its constant shape unless we also mention the way it appears to have moved; this latter change is an essential part of what we’re noting when we remark that the shape of the penny has changed its appearance even as that very same shape appears to have remained unchanged.

The idea, then, is that spatial awareness is essential to visual object perception by being a condition of the possibility of experiencing particular shapes as phenomenally objective, as things that appear to be more than subjective items that happen to be present before the mind. This proposal incorporates a quite different notion of object perception than Campbell’s, and it also differs significantly from Kant’s view of the intuition of space as “that in which alone ... sensations can be posited and ordered in a certain form” (1933: A20/B34): instead, the proposal is that the awareness of space is required not to provide an experiential “container” for visible objects but rather to make for the possibility of objects of visual experience that seem to have more to them than the ways they appear at any given moment. And in contrast to the Apriority Thesis, the idea that spatial awareness is required in this way for visual object perception fits quite well with the available data on RM’s condition: for as we noted in Section 4.2 despite his intact capacities for shape recognition RM usually did not experience objects as able to change their spatial positions; and when they did undergo such changes it was in a haphazard way that made his visual world seem more like an apparition than an awareness of a mind-independent world. Without a visual awareness of how things are situated in space, and how their viewer-relative spatial properties can change under conditions of perceiver- and object-motion, there is no possibility of experiencing the world *as* a mind-independent world at all.

Chapter Two: Vision, Self-Location, and the Phenomenology of the “Point of View”*

1 Introduction

Chapter One concluded by touching on the idea that when you are perceptually aware of the things around you through vision, your experience has a qualitative characteristic that is often called that of presenting its objects “from a point of view”. The experience of looking at a chair from one side is different than that of looking at it from the other, and when you move with respect to an object of visual experience your experience of that object changes in a range of predictable ways. What’s more, and as seems necessary if visual experience is going to play the roles that it must in the motivation and guidance of action,¹ among the things you are visually aware of are the spatial relations that perceived objects bear to your own body. In these ways at least, how things look to you is essentially dependent on where you look at them from.

If we acknowledge all of this while also assuming that there is nothing more to the qualitative character of an experience than the ways in which, due to it, the world appears to the perceiver to be, there is no escaping the conclusion that the location of the perceiver has something important to do with the spatial contents of visual experience.² But what precisely is the nature of this relationship? In particular, does the spatial content of visual experience make manifest only the perceiver-relative locations of “external” objects, or does it include the location *of the perceiver*, as well? It should be clear that there is at least a notional distinction to be made here: for representing egocentric location in a frame of reference that takes the perceiver’s location as its point of origin is not yet representing that point of origin *as* the perceiver’s location; and thus the sensitivity of the spatial contents of visual experience to changes in the location of the perceiver does not entail that those contents include elements which are “self-locating” in the same way as, say, the belief that I am in California or at my desk. The purpose of this chapter is to argue that in the case of human visual experience this gap can sometimes be bridged: in being visually aware of where things are in the world around us we humans are, at least some of the time, visually aware of where in the world we are, too.

Here is how things will proceed. I begin in Section 2 by articulating, and making an initial case for, what I call the *Self-Location Thesis*, that at least some ordinary human visual experiences have explicitly self-locating elements among their perspectival spatial contents. Sections 3 and 4 then develop two different ways of rejecting the Self-Location Thesis, by adopting what I call *Minimal*, *Indexical*, and *Impersonal* views of the contents of visual awareness; but in each case I argue that the alternative strategy in question cannot do full justice to certain aspects of what visual experience is like. Finally, Section 5 considers whether self-locating contents might be an even more ubiquitous feature of visual experience than the initial version of the Self-Location Thesis commits us to.

* Thanks are due to John Campbell, Alva Noë, Sherri Roush, and James Stazicker for comments and discussion. I am especially grateful to Robert Briscoe for some helpful suggestions and also to Tom Avery, whose work on this subject and our exchanges concerning it were instrumental in sharpening my approach to these issues.

¹ The relationships between action and conscious visual awareness are discussed extensively in Chapter Three.

² For more on this idea, see Campbell 2008, which argues that the subject’s point of view ought to be counted among the relata that constitute the character of a visual experience. Note that while the assumption that character bottoms out in content will occasionally be appealed to as a premise in what follows, nothing is meant to hinge on its truth: for further discussion see the second part of Section 3.3 below.

2 The Self-Location Thesis

Suppose I am thinking about the Pythagorean Theorem, and noting this fact I form the further judgment that I am presently engaged in mathematical calculation. Clearly the things I was thinking about gave me a reason for judging as I did, and furthermore they gave me this reason in virtue of their content: if they had been thoughts about something sufficiently different, then I would not have been entitled to judge that it was mathematical calculation rather than, say, recollection of childhood memories that was presently occupying my time. But it would be a mistake to suppose that in order for this to be the case, the fact of my engagement in mathematical calculation must *itself* have been among the contents of my thoughts, at least before I went in for introspective scrutiny and formed the second-order judgment about that calculation: the only things that my thoughts about the Pythagorean Theorem concerned were the relative lengths of the sides of a right triangle, and in noting that aspect of my mental autobiography I introduced into my mental life some new elements that it had not previously contained.

To think about the situation in this way is to treat the relationship between my initial state of mathematical calculation and my second-order judgment about it as one of *representational independence*, where a belief is representationally dependent on a mental state just in case (i) the belief and the mental state have the same content, and (ii) the belief is formed “by taking the mental state ... at face value, in respect of this content” (Peacocke 2000: 264).³ In the case described above, I judge that I am calculating *because* that is what I am doing, and my engaging in that activity *justifies* my judging as I do, but the judgment and the calculation do not have the same content, and so it cannot be that in judging as I do I am simply taking my mathematical thoughts at face value in respect of a content that they and my judgment share. The question whether a thought or other mental state has a given content is thus decided by the question whether there can be judgments with that content that are representationally dependent on that state, i.e. whether there can be judgments with that content that are formed just by taking the content of such a state at face value. That I am thinking about the Pythagorean Theorem is not among the contents of my mathematical thoughts when I am thinking about the Theorem, and thus a judgment as to my thinking takes me beyond the face value contents of my mathematical thoughts themselves.

Another helpful example of representational independence, which Peacocke borrows from Evans (1982: 231), is the judgment, made on the basis of a visual experience of a tree, that *I see a tree*. Clearly the visual experience of a tree is the sort of mental episode that entitles me to judge that I see a tree, by giving me a reason for believing such a thing to be the case: it is because I see the tree that I take myself to see it, and in seeing the tree I am justified in this perceptual self-ascription. Nevertheless, that I see a tree is not, at least usually, among the things *that I visually experience*: the tree itself is an element in the content of my visual perceptual experience, and perhaps – or so we shall argue – there is a way in which I am, too, but my *seeing* is rather the mode in which that content is presented to me, and not an element of the content

³ Actually Peacocke’s definition differs a bit from this, as it only defines representational dependence *simpliciter* for a belief, and not representational dependence with respect to a given experience or other mental state. But for our purposes the account can easily be modified to cut a bit more finely; and indeed this seems worth doing anyway since, as was indicated above, there might be a second-order introspective state that serves as the basis for an autobiographical judgment about one’s thoughts, and it seems that the presence of such a state would qualify the judgment as representationally dependent in Peacocke’s sense.

itself. Seeing a tree is the *kind* of perceptual experience I am presently having, but this is not something that that experience itself manifests as being the case; it is part of the content of my higher-order thought, not of my visual perceptual experience.⁴

Peacocke contrasts to cases like these the example of a self-locating judgment, also made on the basis of the visual experience of a tree, that I am *in front of* a tree. As is the case with my judgments about my thinking and my seeing, in forming this self-locating judgment I am making a claim about myself, and as in the latter case it seems that I am justified in making this claim about myself because of the visual perceptual experience I am having. But can my self-locating judgment, unlike the judgment in which I self-ascribe a visual experience, be formed simply by taking that experience at face value in respect of a self-locating content that it possesses? Peacocke implies that it can be:

Consider ... the everyday case in which an ordinary person forms a belief with the content 'I am in front of a door', and does so for the reason that he sees a door ahead of him. His visual experience represents the door as bearing a certain spatial relation to him. This is so even if he cannot see or otherwise experience his own body on this particular occasion. It would still be true that, taking his experience at face value, he would judge that he is in front of a door. (2000: 264)

So according to Peacocke, *that I am in front of a tree* is a content that my visual experience and my self-locating judgment will both share: simply in virtue of having that experience, I am already in a state that includes my own location among its face value contents. Aside from the tree and its location in perceiver-relative space, *my* location is among the things my experience makes visually manifest, and moreover that experience makes my location manifest *as* my location, thus directly entitling me to the perception-based judgment that that is the location I am in. If Peacocke is right in the way he understands this sort of case, it will supply an affirmative answer to the question that began this chapter, namely that of whether the spatial content of visual experience includes not only the locations of the objects around one, but also makes immediately manifest the location *of the perceiver*, as well. According to Peacocke adequate characterizations of the face value contents of our visual perceptual experiences will sometimes require first-personal indexicals that give the location of the subject.

Moreover, Peacocke's example is especially relevant to the question we began with inasmuch as it is one where neither the perceiver nor his location appears anywhere in the field of view; rather, much as in the way you can have a distinctively visual sense of a three-dimensional object wholly occupying some region of space even as some parts of that object are obscured from view, the idea is that your own location can be an element of your visual experience even when you and your location are strictly *unseen*. This matters if what we are out to understand is whether visual experience has self-locating contents simply in virtue of its "perspectival" character, since this is a distinct question from that of how we should think about those cases in which we simply catch sight of our bodies, whether directly or perhaps in a mirror. And clearly whether we follow Peacocke in treating certain self-locating judgments as representationally dependent on visual experiences will not depend on whether those are experiences in which the self or its location are actually in view: I might judge, on the basis of

⁴ Though note that Harman seriously entertains the possibility that things might be otherwise, that someone's visual experience might present "a tree as seen by her, that is, as an object of her visual experience" (1990: 38). Nothing in what I argue here hinges on whether such a characterization is sometimes or even always correct.

what I see, that I am in front of a tree, even if I have literally no sense of what I or my location presently *look* like – and while having such a sense is surely a plausible requirement on what it takes to see a thing, this is unlikely to affect our opinion of whether such a judgment is representationally dependent. The question whether certain visual experiences include the perceiver's location among their *perspectival* spatial contents is the question whether self-directed judgments can be representationally dependent even on visual experiences that do not bring the perceiver even partly into view.

As a second instance of visual phenomena which seem to supply an affirmative answer to this question, consider J.J. Gibson's discussion of what he calls *visual kinesthesia*, i.e. the way in which visual perception "registers movements of the body just as much as does the muscle-joint-skin system and the inner ear system":

All optical flow vanishes at the horizon and also at the two centers that specify *going toward* and *coming from*. ... Student pilots see where they are going on the basis of this invariant and get better with practice. Drivers of cars see where they are going if they are paying attention. Viewers of a Cinerama screen see where they are going in the represented environment. A bee that lands upon a flower must see where it is going. And all of them at the same time *see the layout of the environment through which they are going*. ... Vision picks up movements of the whole body relative to the ground and movement of a member of a body relative to the whole. Visual kinesthesia goes along with muscular kinesthesia. The doctrine that vision is exteroceptive, that it obtains "external" information only, is simply false. Vision obtains information about *both* the environment *and* the self. (1986: 182-83)

Moreover, Gibson continues, the experiential counterpart to the visual perception of one's own movement is the similarly visual perception of one's rest:

Vision, of course, is also *statethetic*, if one wants to be precise about words, in that it picks up *nonmovement* of the body and its members. But since nonmovement is actually only a limiting case of movement, the term *kinesthesia* will do for both. The point is that a flowing and an arrested optic array specify respectively an observer in locomotion and an observer at rest, relative to a fixed environment. Motion and rest are in fact what an observer experiences with flow and nonflow of the array. (1986: 183)

There is something clearly relevant to our purposes in what Gibson is claiming here, even if for the reasons canvassed in Section 2 of the Introduction his emphasis on the notion of the "optic array" leaves his emphasis at a bit of a distance from the question we are considering. In our terms, the crucial moral is that at least in some of the cases where a perceiver moves through space *his own motion or rest* can be among the things of which he is visually aware: he can experience the movement of his body *as* his movement or his motion *as* his rest, and not in virtue of having any part of his body in view. If this analysis is correct, then it is possible for ascriptions of self-movement or -nonmovement to be representationally dependent on visual experiences in the same kinds of ways that Peacocke argued self-locating judgments can be: the judgment that I am stationary, or that I am moving (perhaps: at a certain rate and in such-and-such a direction), need not move me beyond what is made immediately manifest in visual experience, but rather can take those contents simply at face value. A perceiver's judgment that

he is in motion can just be a way of giving one's endorsement to the way visual experience reveals things to be.

When we reflect more closely on the matter, we find good reason to think that the visual experience of self-location and self-motion and -rest are closely interwoven, such that if experience of one of these sorts is possible then we will find the corresponding possibility, too. For part of what it is for a spatial content to be self-locating rather than merely perceiver-relative is for it to represent the location of the perceiver as potentially on a par with any other location, as one among many possible locations in an objective spatial world. It seems impossible to do this, however, without also being able to represent changes in that location in a situation where other represented objects remain unmoved; and correspondingly, it seems impossible to represent self-motion in this sort of way unless the location of the subject can be given an objective content as well. So evidence that the face value contents of visual experience can include first-personal elements of one of these sorts will be evidence that it can include the other sort of first-personal content, too.

We can title the idea that visual experiences can include among their face value perspectival spatial contents these sorts of aspects of the location, rest, and motion of the perceiver the *Self-Location Thesis*:

(SLT) At least some visual experiences have *self-locating contents*: that is, their face value perspectival contents include spatial properties of the perceiving subject, such that in judging as to his or her own location or motion or rest a perceiver can simply be taking the contents of such experiences at face value in virtue of contents that they and the corresponding judgments share.

Motivated in the way we just have, the Self-Location Thesis should seem intuitively obvious and perhaps even incontrovertible; but as often happens we will discover quite quickly that it is philosophically contentious, and it will take further argument to show that it is correct.

3 The Minimal View

According to the Self-Location Thesis, it is possible for judgments of one's own location and motion or rest to be representationally dependent on ordinary visual experiences: that is to say, it is possible to self-ascribe certain spatial properties simply by taking the perspectival spatial contents of one's visual experience at face value; and correspondingly, it will sometimes be impossible to specify the contents of such an experience without mentioning the apparent spatial properties of the perceiver. Thus Quassim Cassam:

... the spatial content of perception is egocentric. As Husserl puts it, 'all spatial being necessarily appears in such a way that it appears either nearer or farther, above or below, right or left' (1989: 166). The important point is that in egocentric spatial perception the objects of perception are experienced as standing in such spatial relations *to the perceiver*; in Husserl's words, 'the "far" is far from me, from my Body' (ibid.). ... Egocentric spatial perception can therefore be described as self-locating; in experiencing objects as spatially related to one, one literally experiences the bodily self as located in the perceived world. (1994: 52-53)

This thesis has recently been challenged by John Campbell, who draws a crucial distinction to which Cassam may be being insufficiently sensitive, namely that between *mere* egocentric spatial content and spatial content in which the perceiver's location is explicitly represented:

The notions 'above', 'below', 'right', 'left', 'in front', and 'behind' are usually called 'egocentric' notions, the idea being that they define positions with respect to the subject: the perceiver and agent is taken to be the origin of the frame of reference, and places are identified by their relations to that subject. But there is a basic distinction that we have to draw here between what I shall call relational and what I shall call monadic egocentric spatial notions. Relational egocentric notions are those that we use when we say, for example, 'He is sitting on my left', 'The chasm yawned before him', 'Look behind you', and so on. These notions specify the person whose right or left, up or down is in question. They are two-place notions: 'x is to y's left', 'x is below y', and so on. Now in stating the spatial content of vision, we do not seem to need these relational notions. We do not need the general conception of something's being to the right or left of an arbitrary subject. Rather, we need the more primitive monadic egocentric terms. These are notions such as 'x is to the right', 'x is below', and so on. An animal could quite well have spatial vision even though it did not have the relational egocentric notions; it could not represent anyone else's left or right, only its own. But it is not even as if its vision makes explicit the spatial relations that things bear to it ... Its vision represents things as 'to the right' or 'above'; it does not seem correct to say that it represents things as 'to my right' using the relational notion, because of the lack of generality in whose left or right can be represented. And the same seems to be true of ordinary human vision. It represents things as 'to the right' or 'above' using the monadic egocentric notions, rather than the relational terms. (2002: 184)

We will have more to say in short order about the details of how this contrary proposal can best be understood, but for now it is important to note that nothing in what Campbell says here commits him to denying that an ordinary human perceiver can form representationally *independent* judgments of self-location on the basis of his visual experience, nor need he insist that there is no sense at all in which such a perceiver can be said, *à la* Gibson in the passages quoted above, to *see that* he is, say, in motion or at rest, or located in such-and-such a place. Rather, Campbell can treat these latter locutions in the way that we ordinarily think of, say, someone's being able to see that (or: have it look to him as if) his house has been burglarized even though its having been burglarized, as opposed just to there being a broken window and footprints in the flowerbeds, is not a face value content of his visual experience, but is rather something he is immediately inclined to take to be so on the basis of how things visually appear. I will argue in this section, however, that such an alternative position falters on the question of what the non-self-locating face value contents of the visual experiences in question might be: there is simply no way to specify what is special to the contents of certain kinds of visual experiences without making explicit mention of the visually apparent spatial properties of the self.

3.1 Content and Concern

According to the view Campbell expresses in the passage quoted above, we can account for the perspectival spatial content of visual experience by identifying the spatial locations of perceived objects solely using monadic egocentric terms like “ahead” and “to the left”, predicates which make no mention of the perceiver’s location. Such a monadic spatial vocabulary is, of course, clearly *indexed* to that location: it is because the perceiver’s body serves as the origin of the egocentric spatial reference frame that the egocentric concepts have the contents that they do in any given instance, and that those contents will therefore change when there are appropriate changes in the spatial relations between the perceiver and the objects around him. How exactly can we make sense of the idea of spatial contents that are, as we can put it, *self-indexed* in this way without being fully first-personal in the manner of familiar indexical concepts like “I”, “me”, and “my”?

Imagine an organism that is visually sensitive to the egocentric locations of things but nevertheless lacks the capacity to represent itself under first-personal modes of presentation: the organism will act differently depending on whether an object is to its right or its left, and these actions will be a direct response to the location that the organism’s visual perceptual states represent the object as standing in with respect to it. Clearly it is possible for our imagined organism to be visually sensitive to the perceiver-relative locations of things *only* in this way, such that it behaves differently depending on whether a perceived object is to its right or its left but it is nevertheless wrong to say, as it were from the organism’s perspective, that any of those states makes manifest the fact that, say, an object is to the left *of me*: rather, the object is simply presented as having a certain location, and the identity of that location is individuated by its spatial relations to the perceiver even though the perceiver’s location is never represented as such. (This is likely to be the right way to think about the self-specifying information picked up by the visual system of Gibson’s flying bee, which suggests that in this respect his account of visual kinesthesia is a bit too coarse-grained.) As Ruth Millikan puts it, in a case like this we have a perceptual state that is “‘ego-centered’ in that the ego is so central to it that it doesn’t even need to be mentioned” (2002: 179): the location of the perceiving subject is “implicit” in the state’s spatial content without being explicitly included therein.

It is important that we not confuse the idea of a system of spatial representation that is egocentric in this sense with that of one that treats any change in perceiver-relative location as a change in location simpliciter.⁵ For example, Acredolo 1990 describes an experimental paradigm in which an infant is seated inside a room with two windows, at one of which a friendly face consistently appears following an auditory signal. After several preparatory trials the infant is rotated 180° to face the opposite direction, the tone is played, and we wait to see whether the expectant infant orients her head to the same side of the room where the face had previously appeared: if she does, then this is evidence that she had individuated the window’s location in virtue of something other than its relations to her own body; while if not, then location within a body-centered frame of reference is taken to exhaust the infant’s grasp of space. But clearly it is possible for the infant to recognize that what was to the left a moment ago is now to the right even without representing herself as now occupying a different location, or as having moved in between: all that is required for success in the Acredolo paradigm is that the system of spatial representation whose contents we capture in monadic egocentric terms not *exhaust* the subject’s

⁵ Here I am especially indebted to correspondence with Robert Briscoe.

means of identifying sameness or difference of location; there needs to be a representation of location in “objective” space to complement that of the relations of things to the subject’s own body. And while this further dimension of spatial representation will have to draw on a perceptual sensitivity to changes in the subject’s location, this does not yet mean that it will be appropriate to specify its contents in explicitly self-locating terms. At least in principle, our hypothetical organism can experience a stable world without being able to experience itself as located anywhere within it.

Perhaps the most detailed account of what differentiates a merely self-indexed system of spatial representation from one possessing genuinely self-locating contents is developed in John Perry’s “Thought Without Representation”. Perry begins from the same thought that opened Section 1 above, that visual experience reveals where things are located with respect to the perceiver; yet like Campbell he maintains that we should acknowledge this without committing ourselves to drawing on explicitly first-personal language in giving an account of vision’s contents:

I see a cup of coffee in front of me. I reach out, pick it up, and drink from it. I must then have learned how far the cup was *from me*, and in what direction, for it is the position of the cup relative to me, and not its absolute position, that determines how I need to move my arm. But how can this be? I am not in the field of vision: no component of my visual experience is a perception of me. How then can this experience provide me with information about how objects are related to *me*? (1993: 205)

Perry proposes that we can make sense of this possibility by drawing an analogy with the thought and verbal behavior of an imagined linguistic community whose members are sensitive to what happens around them yet do not conceive of their location as one among many possible others. For an ordinary English speaker, Perry explains, an utterance like “It is raining” expresses a proposition whose content includes the location from which it is uttered: we cannot state what is meant by such a statement without including in our expression of its content a term referring to the place where it is spoken. But it is possible for an utterance’s location to bear on its content in ways other than this, as in the following example:

Consider a small isolated group, living in a place we call Z-land. Z-landers do not travel to, or communicate with, residents of other places, and they have no name for Z-land. When a Z-lander sees rain, he will say to others not in a position to look outdoors, *It is raining*. His listeners then act appropriately to there being rain in Z-land: they close the windows in Z-land, cancel plans for Z-land picnics, and grab umbrellas before going into the Z-land out-of-doors. They have no other use for “It is raining.” They do not call their sons in far-off places, or listen to the weather news, or read newspapers with national weather reports. (Perry 1993: 212)

In this case, like that of a community of ordinary English speakers, the fact that the Z-landers are in Z-land rather than somewhere else is clearly relevant to how we should understand what they say: when a Z-lander says “It is raining” his utterance is true if and only if it is raining in Z-land; and we could just as well imagine, for example, a group of speakers who say “It is raining” only when it is raining somewhere other than where they are, in which case those utterances would be caught up quite differently in the speakers’ and hearers’ forms of life. (Perhaps they only care about the weather in Hawaii.) But in contrast to ordinary English speakers, it seems improper to

say of the Z-landers that they *mean* anything about Z-land when they say “It is raining”, since as things stand it’s impossible for them to use the predicate “... is raining” to make a claim about the weather anywhere else: as Perry suggests, whereas ordinary English utterances of such a sentence pick out a semantically complex relation involving both times and places, a Z-lander semanticist should treat “is raining” as a simple predicate that holds only of times, since Z-landers don’t understand, and don’t have the linguistic resources to express, what it would be for it to be raining anywhere other than where they are. (Indeed: as we are imagining the Z-landers they do not understand what it would be to *be* anywhere other than where they are; they simply don’t have the concept of “where I (/we) are” as an objective *place* at all.)

Thus Perry proposes that we mark off this contrast by saying that while ordinary English utterances of “It is raining” are (partly) *about* the speakers’ putatively wet and spoken-from locations, such utterances in the mouths of Z-landers merely *concern* the place they are spoken from. It is because of this that the Z-landers’ rain-thoughts and -sentences can bear on their behaviors in the ways that they do by, say, leading them to cancel plans or get an umbrella before going outside rather than (as in the case of the imagined subjects for whom talk of rain always concerns the weather someplace else) calling faraway family members to encourage them to do the same, even though the Z-landers don’t actually conceive of Z-land as one place among many. And on Perry’s account, we can extend this idea to the case of visual experience by saying that while visual experiences contain “information *concerning* ourselves”, and thus are able to undergird the necessary connections between perception and action, there is “no need for a self-referring component ..., no need for an idea or representation of ourselves” (1993: 219), for these sorts of connections to be possible.

Here, then, is how we should think about the first alternative to the Self-Location Thesis: not as a denial of the obvious datum that the contents of visual experience reflect, and perhaps are even rooted in a system of spatial representation that is sensitive to, the location of the perceiving subject; but rather as the claim that explicitly self-locating elements are never among the immediate contents of visual experience. Hence on such a view, when I come to judge, on the basis of what I see, that a tree is five yards away from me, it cannot be because I am taking my visual experience simply at face value. For *my* location, unlike that of the tree, is not among the things I experience in such an episode of visual awareness; rather, it serves only as an implicit index that determines its egocentric spatial contents.

Call the view endorsed by Campbell and Perry the *Minimal View*. According to proponents of the Minimal View, an adequate description of the face value contents of visual experience need never use explicitly first-personal terms like “I” or “me” in giving its perspectival spatial contents; rather, the location of the perceiver serves only as an origin point for the frame of reference that specifies the represented locations of perceived objects in an egocentric space. Proponents of the Self-Location Thesis, by contrast, deny the Minimal View and hold instead that judgments of one’s own location and motion or stationariness can be a way of taking the contents of visual experience at face value: certain of the perceiver’s spatial properties can a basic aspect of the way things look to him. The Minimal View thus challenges the Self-Location Theorist’s claim that visual experience is ever explicitly self-locating, as opposed to merely perspectival in virtue of having egocentric spatial contexts that are indexed to the location of the perceiver.

3.2 Is the Minimal View Phenomenologically Adequate?

If we restrict our attention to the sorts of visual experiences enjoyed by an exclusively stationary observer, it is hard to deny that the Minimal View has considerable appeal; but as we saw in Section 2 it is not just those experiences whose perspectival spatial contents the Self-Location Thesis has been invoked to explain, and the Minimal View looks considerably worse off when we apply it to those experiences that Gibson labels as visually “kinesthetic”. As an example of such an experience, consider Stephen Palmer’s description of the visual illusion of vection induced by placing a subject inside a rotating cylindrical drum (see Figure 2 below):

... if you were seated inside a large, opaque, cylindrical drum with vertical stripes painted on it, ... and if the drum were rotating, you would soon perceive the drum as stationary and yourself as spinning in the opposite direction inside it. This experience of self-rotation is so compelling that many people become dizzy and nauseous [*sic*], very much as they would if they were absolutely rotating. In fact, however, they are quite stationary; only the cylinder around them is moving. (1999: 505)

Note that while this is a case where an illusory sense of self-motion is induced purely by optic stimulation, there are any number of more ordinary cases, including ones where the character of visual experience is arguably impacted by non-visual sources, that could have served our purposes just as well: as Palmer suggests (1999: 505), the ordinary veridical experience of driving a car through a stationary environment seems to be one in which your visual sense of motion is affected not just by the light that strikes your eyes, but also by processing in your vestibular system in your inner ear. But for present purposes, Palmer’s example of an illusion of vection provides a helpfully simplified case that can reveal where the Minimal View comes up short.




Figure 2: Visual illusion of self-motion (vection) induced by a rotating drum. The cylinder rotates counterclockwise (solid line), and a stationary observer placed inside it soon has an illusory visual experience of himself as rotating clockwise (dotted line) and the drum as stationary. Based on Figure 10.3.2 from Palmer 1999: 505.

So you are in the drum, and it is spinning to the left, but it seems to you as if you're spinning to the right – we're going to be referring back to this example quite a lot, and so to avoid begging any questions about the best way to describe this experience let's simply call it "E₁". According to the Self-Location Thesis, that we offer such a description of how things seem to the subject in E₁ derives from the fact that in experiences like this one, the perceiver's motion is among the face value contents of visual experience: it is *yourself* whom you perceive (or misperceive) as spinning, while the drum around you visually appears to you to be stationary. And it is exactly such explicitly self-locating elements that the Minimal View holds never to be among the face value contents of visual experience.

But is there a different way for a defender of the Minimal View to characterize the face value contents of visual experiences like E₁? If we are willing to treat an experience of motion as a series of static visual "snapshots", each of which reveals the perceiver-relative positions of things in an egocentric space that is self-indexed but never explicitly self-locating, it appears that Campbell's proposal will do quite well: supposing that the cylinder is spinning counterclockwise, a line inside it that appears *to the right* at t₁ will appear *straight ahead* at t₂, *to the left* at t₃, and so on. But the problem with this response is that these very same egocentric spatial contents will also be possessed by an experience, which we can call "E₂", in which it is the *cylinder* that appears to be spinning, while the subject has no sense of self-motion at all; and as Palmer notes (1999: 505), a person placed inside a rotating drum will undergo a sort of gestalt shift from a briefly veridical experience of the motion of the cylinder to an illusory experience like E₁. If we respect the constraint that aspects of qualitative character reduce to aspects of face value content, the "snapshot" version of the Minimal View has so far failed to state the dimensions along which experiences like E₁ and E₂ differ.

Clearly one further thing that the Minimal View allows us to say of Palmer's example is that only once the illusory experience E₁ has set in does the *cylinder* appear to be stationary, while in the short period before the illusion there will be a veridical experience of its moving from right to left. This is certainly *a* difference between the contents of E₁ and E₂, but does it exhaust the ways in which their characters diverge? It seems not: for in addition to this difference there is an *apparent motion* present in the illusory experience E₁ that does not have a counterpart in E₂, as well as an *apparent motionlessness* in the latter that is absent from the former; and it seems to be this that the subject is making note of when he says of the illusory experience that not only does the cylinder appear to be motionless, but also that *he* appears to be spinning, and of the veridical experience that the cylinder appears to be spinning while *he* is apparently still. When we recall from our own past experience what an illusion of vection is like, the phenomenological accuracy of such descriptions is simply undeniable: what it is like to have an IMAX screen or a moving truck or train car next to one induce an illusion of forward or backward motion is visually quite different from what it is like to have a mere experience of motion in one's environment, say of the neighboring train car moving forward or the images on the IMAX screen simply shifting around. Thus unlike the example of the seeming burglary, we simply misdescribe what visual experience is like if we insist that in cases like these judgments of apparent self-movement are drawn by inference from visual experiences whose face value contents are entirely non-self-locating; rather, these subjective reports of the apparent presence or absence of self-motion mark aspects of visual experiences every bit as basic to their contents as the apparent motion or rest of the things one has in view. It is impossible to say all that needs to be said about the differing qualitative characters of E₁ and E₂ without attributing to the former an explicitly first-personal content involving the changing location of the perceiver.

3.3 Two Possible Responses

How might a defender of the Minimal View respond to this challenge? One possibility is to push back against the demand for an account of the qualitative difference between E_1 and E_2 by denying that the experiential relevant difference is a *visual* one: the idea would be that things are *visually* the same whether I am moving in one direction or the things around me are all moving in the opposite one, and so it is only in some non-visual modality – like nausea or dizziness, say – that any experiential differences show up.⁶ This line of response might be thought to find motivation in the fact that, as Palmer's example and any number of everyday examples of visually induced illusions of self-motion make clear, self-motion and world motion can easily be mistaken for with one another because of the similarity of optic stimulation: isn't this just because the corresponding visual experiences are simply the same?

The fundamental problem with this way of defending the Minimal View against our objection is that it runs up against a number of basic intuitions about what it takes for a dimension of experience to count as visual. For example, when you look at a visually ambiguous figure like the Necker cube and its appearance undergoes a shift of aspect from having one apparently protruding side to having another, there is no sense in denying that despite the sameness of optic stimulation it is your *visual* experience that undergoes a qualitative change: perhaps you are also inclined to do extra-visual things like reach and grasp the figure differently or form different judgments about it, but these differences seem to be a *product* of the fact that there is a distinctive change in how you visually experience the figure to look. And what Palmer's and Gibson's example help to bring out is that there can be a similarly visual difference between an illusion of vection like E_1 and the veridical experience of corresponding motion in the surrounding world, a difference it seems impossible to capture without describing how things look in explicitly self-locating terms. Aside from a stubborn unwillingness to abandon the Minimal View, it is hard to see what could motivate someone to insist that the difference between E_1 and E_2 must be treated so differently than ordinary gestalt shifts in the visually apparent structure of an optically constant stimulus.

If it is agreed that the difference between E_1 and E_2 is at least in part a genuinely visual one, and that the Minimal View cannot account for this difference in the terms proposed in Section 3.1, then the only remaining option short of significantly modifying the Minimal View is to bite the bullet and jettison the assumption we took on at the start, namely that differences between the qualitative characters of visual experiences are reducible to corresponding differences in their contents. For as we advanced our initial argument against the Minimal View in Section 3.2 it was precisely in striving to meet this requirement that we were forced to ascribe self-locating contents to E_1 and E_2 : the idea was that the qualitative characters of the respective visual experiences of self- and world-motion had to be grounded in differences in content, and that the only kinds of contents that seemed able to do this job were explicitly self-locating ones. But what if this demand for the reduction of character to content is rejected?

The most natural place to turn to motivate such a response is to Christopher Peacocke's *Sense and Content*, which argues as follows for the existence of non-representational properties in visual experience:

⁶ This objection has been pressed on me by Sherri Roush, in conversation, though the way I develop it here is my own.

Suppose you are standing on a road which stretches from you in a straight line to the horizon. There are two trees at the roadside, one a hundred yards from you, the other two hundred. Your experience represents these objects as being of the same physical height and other dimensions; that is, taking your experience at face value you would judge that the trees are roughly the same physical size ... Yet there is also some sense in which the nearer tree occupies more of your visual field than the more distant tree. This is as much a feature of your experience itself as its representing the trees as being the same height. The experience can possess this feature without your having any concept of the feature or of the visual field: you simply enjoy an experience which has the feature. It is a feature which makes Rock say that the greater size of the retinal image of the nearer tree is not without some reflection in consciousness, and may be what earlier writers such as Ward meant when they wrote of differences in extensity. It presents an initial challenge to the Adequacy Thesis [i.e., the view that we can adequately characterize an experience simply by giving its content], since no veridical experience can represent one tree as larger than another and also as the same size as the other. The challenge to the extreme perceptual theorist is to account for these facts about size in the visual field without abandoning the AT. (1983: 12)

According to Peacocke, there are dimensions along which the difference of a faraway tree differs from that of a nearby tree of the same size that cannot be accounted for by differences between those experiences' respective contents: that the one tree is apparently farther away than the other is *part* of what makes them different, but there is a further experiential difference that outstrips this difference in content and that is determined by the spaces that things take up in the visual field. There are aspects of visual experience that go beyond the way it presents the world to be.

A defender of the Minimal View, then, might object to our arguments for the Self-Location Thesis along similar lines: "You have said that changes in character are impossible without changes in content, and so that the only way to account for the particular experience of *motion* that we find in an illusion of vection like E_1 but not a corresponding experience of world-motion like E_2 is to say that only in the former case does the subject have a visual experience of *his* motion in an otherwise stationary environment. But this overlooks the fact that visual experience has purely qualitative features that go beyond the ways it represents things to be, a feature corresponding to the regions that objects occupy in the subject's visual field: thus in an experience like E_1 there is, in addition to the apparent motionlessness of the surrounding environment, also the leftward visual field movement of the lines inside the cylinder. This is the feature of experience that makes Rock say that despite the visual experience of objects' constant positions across movements of the perceiver, 'this angular motion [of objects with respect to a frame of reference defined by the subject's eyes] is in some sense perceived because we are aware that objects are changing their location in the field of view' (1983: 257). And it is only this feature that you are pointing out when you note that there is visually apparent motion in an illusion of vection; we need not treat this experience of motion as a case of changing visual experiential contents."

So the defender of the Minimal View argues that the difference between an illusion of vection and the corresponding experience of a moving world is not a difference in the visually apparent location of the perceiver, but rather a purely phenomenal difference deriving from the changing positions of things in the visual field. But as soon as we state the response in this way, it becomes clear that rejecting the dependence of character on content does not advance the Minimal View even one step beyond the position we already rejected in Section 3.2: for these

allegedly “purely qualitative” changes in the character of visual experience are, *just like* the changing positions of things in an egocentric space with contents defined by monadic predicates indexed to the perceiver’s body, clearly *shared* by experiences like E_1 and E_2 alike. In *each* of these cases, whether it is the surrounding environment or (as the Self-Location Thesis would have it) oneself that appears to be moving, the environing objects will be changing their positions in the subject’s field of vision, and so according to a theory that postulates a purely qualitative dimension of experience of the sort in question *each* case will therefore involve a corresponding experience of leftward motion across the phenomenal visual field. Since this qualitative characteristic is something that the two experiences have in common, it cannot be used to explain how there is an aspect of apparent motion in thevection illusion that is missing from the experience of a moving world. The premise connecting content and character is a plausible and dialectically helpful assumption, but it can be discharged without doing any harm to our case against the Minimal View.

4 The Indexical and Impersonal Views

According to the Minimal View, the perspectival character of visual experience arises from the fact that its spatial contents are defined, at least in part, by a reference frame that is indexed to the body of the perceiving subject; but we can state what those contents are without referring to the perceiver’s location as such. As we saw, such a proposal comes up short when it encounters what Gibson calls the phenomena of visual kinesthesia and statesthesia, for in such a case there can be visual experiences of motion or rest that is absent from corresponding experiences of the movement or stationariness of objects in the surrounding environment, while the only ways that the Minimal View allows us to describe changes and constancies in perceiver-relative position apply equally to each member of this pair. Thus we concluded that the only way to account for the ways that such experiences differ is to attribute to an experience like E_1 face value contents involving the perceiver’s visually apparent location. But this need not be the end of philosophical resistance to the Self-Location Thesis, as there are two further conceptions of the content of visual experience that remain to be considered.

On the one hand, according to what we can call an *Impersonal View* of visual self-awareness, there is a crucial distinction to be drawn between the claim that (1) a visual experience like E_1 includes among its face value contents the apparent motion of *something* that is not experienced as moving in corresponding cases of apparent world-motion from the claim that (2) it is *oneself* who, in a case like E_1 , is visually experienced as having such a property. (As an approximate comparison, think of catching sight of yourself in a television monitor but not realizing who it is that you’re perceiving: your visual experience makes you aware of what is *in fact* yourself, and what is *in fact* your location, but they are not presented to you under those specific modes.) Put somewhat differently, the strategy here is to allow, as (2) would have it, that the locations of our bodies can be face value aspects of the perspectival spatial contents of visual experience, but deny that (as in (1)) these are cases where we visually experience our locations *as* our locations: there can be a bit more to the face value content of visual experience than the Minimal View allows, but nevertheless that content is never appropriately expressed in explicitly self-locating terms.

On the other hand, an *Indexical View* of the phenomena in question holds that visual kinesthesia and statesthesia can be understood in terms of semantically primitive experiential contents that do not involve the motion or rest of any *thing* at all: on this account, the experience

of one's own motion or rest is akin to the experience of pain or emotion, at least on some philosophical understandings of what those experiences involve; the idea is that the experiential contents in question are "causally indexical", in that their significance to the perceiver is exhausted by the suite of behaviors to which they ordinarily give rise. This notion of causal indexicality is due to Campbell:

Unstructured uses of 'is heavy' and 'is hot' may relate to the causal impact of the thing upon the subject, rather than being uses of some observer-independent system of classification. Or again, a notion such as being within reach seems to have immediate implications for the subject's actions. The most immediate effect of judgements made using this notion is that the subject will try to contact things within reach but will not try to contact things judged to be out of reach. This predicate is not a first-person one [like 'is within reach of *me*']. A creature could use representations of things as within reach or out of reach without having the ability to think using the first person. (1993: 44)

As Campbell's formulations suggest, the Indexical View is essentially a variant of the Minimal View, and not a genuine alternative to it; but it will be helpful to consider it under a separate heading, as it addresses the phenomenology of self-movement more explicitly and is subject to a different set of philosophical problems than the Minimal View as we initially articulated it. Only if the Indexical and Impersonal Views are defeated are we entitled to assert the truth of the Self-Location Thesis.

4.1 Against the Indexical View

The best way to get a handle on the Indexical View is to think about how it might be put to work in our thinking about the perceptual states of an organism with an evident visual sensitivity to the difference between self- and world-motion but which, for one reason or another, we are disinclined to treat as the subject of genuinely self-locating visual experiential states. Such an organism will, for example, adjust its gait or posture so as to maintain balance in cases of (real or illusory) visual kinesthesia, while in corresponding cases of evident world-motion its behavior will be quite different; nevertheless, we can explain these behavioral differences in the same sort of way as Perry understands the rain-behavior of his imagined Z-landers, by supposing that just as they can be behaviorally sensitive to the weather around them without being able to think in explicitly self-locating terms, so the organism's visual perceptual states can have contents that *concern* its own motion and rest, without explicitly representing the organism's own position as such.⁷

Can we, however, make sense of the idea that *human* visual experiential states have contents that are impoverished in this way? As Campbell notes, there are many instances in which we can easily come up with ways to reformulate first-personal predicates like "is a weight I can easily lift" or "is too hot for me to handle" in causal indexical terms ("is heavy" and "is hot" respectively: see Campbell 1993: 45); and a similar thing holds true of monadic spatial predicates like "up ahead" or "to the right", which are as much a part of our ordinary spatial vocabulary as self-relational terms like "ahead and to the right of me". By contrast, there is no straightforward way for us to express the contents of experiences like E_1 and E_2 in the way that

⁷ Many thanks to John Campbell for instigating the discussion of the following paragraphs.

the Indexical View requires us to: perhaps the closest we can get is to use cooked-up sentences like “It is moving” and “It is at rest”, where these are somehow akin to the English “It is raining”, which simply reports the presence of rain without ascribing a property to any object (as in “It is red (or heavy, or ...)”). Yet even this analogy is imperfect, as the English predicate “is raining” can be applied to places other than where one is, whereas the imagined causally indexical terms for self-motion and -rest are essentially specific to oneself. If we think that it counts against an account of the ordinary contents of visual experience that it leaves us with no natural way even to state what those contents *are*, then this aspect of the Indexical View will give us real cause for concern.

Moreover, this same aspect of the Indexical View commits its adherents to rejecting the possibility of integrating the contents of visual experience with our everyday thought about ourselves and the world around us. For even if it makes sense to think that there are cognitively significant aspects of the content of experience that cannot be given a verbal expression, the Indexical View treats the self-“concerning” contents that differentiate experiences like E_1 and E_2 as involving fundamentally different features from those involved in the experience of the motion and rest of ordinary physical objects. Treating the contents of visual experience in this way leaves it a mystery how such experience could put us in a position to think of self-movement as on a par with the movement of objects other than oneself, and so of oneself as an ordinary physical object among many others. Clearly this is not much of an objection when the perceivers whose experiential states we are trying to understand are organisms who *can’t* think of themselves in this way, but when it comes to organisms with a higher degree of cognitive sophistication the Indexical View simply seems unmotivated; human visual experience is not disconnected in this way from our understanding of ourselves as located in an objective spatial world.

But the most fundamental problem with the Indexical View has to do not with the relationships between experience and thought or verbal report, but rather with the experiential *similarity* between the visual experience of one’s own motion or rest and the motion and rest of the objects around one. There is, of course, something experientially distinctive about the visual awareness of one’s own motion or rest “from the inside”, as opposed to that of the motion or rest of an object that is present somewhere in one’s field of view. Yet when we stated what was distinctive about an experience like E_1 that made it different from a corresponding experience of world-movement like E_2 we found ourselves forced to characterize the fundamental difference precisely in terms of the presence or absence of *motion and motionlessness*: there was no other way of stating the difference that seemed able to do the trick. The Indexical View, however, attempts to ground the distinctive aspects of experiences like E_1 and E_2 contents that do not involve the *properties* of motion or rest at all, as part of what it is to be a property is to be able to be instantiated by a range of different objects, whereas the causally indexical features invoked by the Indexical View to account for the differences between experiences like E_1 and E_2 are essentially self-concerning, and thus are not properties of any objects at all. That there can be organisms whose visual perceptual contents are impoverished in this way is beyond dispute, but phenomenological reflection on what visual experience is like in our own case indicates that this is precisely one of the ways in which their perceptual states differ from ours.

4.2 Against the Impersonal View

While the Indexical View treats the experiences of motion in an experience like E_1 and of motionlessness in E_2 as involving nothing more than the experience of semantically unstructured properties that are not applied to any subject at all, the Impersonal View allows that in cases like these one visually experiences the motion and rest of *something*, but denies that that “something” is visually experienced as the perceiver himself. This is a way of rejecting the Self-Location Thesis because it rejects the claim that in stating the face value contents of visual experience we need to use explicitly first-personal language, and so treats the relationship between, say, my experience of a tree and the judgment that I am in front of a tree as one of representational independence: perhaps there is less of a gap to be bridged than the Minimal View supposes, but the visual experience and the visually-based judgment differ in content nonetheless. But the challenge for adherents of the Impersonal View, however, is to specify what, if they are not explicitly self-locating, the relevant contents of visual experience could be.

To see why this question indicates a significant hurdle for the Impersonal View, consider a response to it that would obviously be inadequate, namely that what differentiates experiences like E_1 and E_2 is that an experience of the first sort includes the content that *someone* is moving, while the latter does not. Marking the visual difference between these experiences in this way is clearly not a concession to the Self-Location Thesis, but it is equally clearly not an adequate account of how the contents of these experiences diverge: for example, the experience of an otherwise stationary environment in which a person *other than* oneself is moving is also a case in which “someone” visually appears to be moving, but the way in which experiences like E_1 and E_2 differ is more specific than this. The challenge for a defender of the Impersonal View is to mark this phenomenal difference in a way that accounts for its specialness – i.e. the way in which the sort of visual experience that accompanies one’s own movement is phenomenally different from that of an ordinary object – without using explicitly first-personal language in saying what it is.

Given this challenge, another version of the Impersonal View which might seem more promising is to say of the contents of E_1 that in such an experience it is *this person*, or perhaps *this body*, that appears to be in motion in an otherwise stationary environment. The hope here would be that the perceptual demonstrative ‘*this*’ marks the experiential distinctiveness of the person (or body) that visually appears to move, thus differentiating such an experience from a corresponding one in which the thing in apparent motion is simply an object in the field of view. But any seeming promise of this proposal is revealed to be illusory when we push it a bit further: for of course the case in which an object in the surrounding world is perceived to be in motion can *also* be described by using perceptual demonstratives to refer to that individual, and even if the demonstratives are being used differently in these cases (i.e., in that they are referring both to different individuals and – perhaps more importantly – to individuals occupying different locations in perceiver-relative space) that difference is not “exhibited”, as it were, simply by citing demonstrative contents in this very generic way. What is needed is a possible visual experiential content that is non-self-locating but nevertheless subjectively proprietary, in the right sort of way, to the thing that occupies the perceiver’s location.

On reflection it seems that the only way for advocates of the Impersonal View to meet this demand is to characterize the perspectival spatial contents of visual experience partly in terms of the apparent location of something like *the origin of the visual reference frame*, or

perhaps just *the point of view* itself.⁸ On this way of developing the Impersonal View, what a proponent of the Self-Location Thesis wishes to describe as a case where it is *myself* whom I experience as moving or at rest, or as located at such-and-such a point in space, is better treated one of the apparent motion, rest, or location of what Husserl (1989: 166) calls the “*zero point*” of the egocentric spatial orientations. Much as I come to believe, in a representationally independent way, that my house has been burglarized because of the muddy footprints and the broken window that I see, so my judgment that I am, say, in motion is based on a visual experience that does not possess that particular content, but rather manifests the changing position of something that is not itself me, but whose location is essentially the same as my own. Strictly speaking, the contents of an experience like E_1 are silent when it comes to the spatial properties of the perceiver, but they do involve the apparent motion of the visual point of view, together with constant locations and changing egocentric positions of the visible objects that surround it.

But is it really plausible that concepts like those of visual reference frames or points of view should be drawn on in giving the face value contents of ordinary episodes of visual awareness? One natural reason to resist this idea would be if we held that it is impossible for a subject to have an experience with a certain content without possessing, at least being in a position to possess, the corresponding concept: for clearly it seems possible for prelinguistic children and at least some non-human animals to have visual experiences like E_1 , even as the concepts of a point of view or the origin of a reference frame are well beyond their reach. But in fact we can identify the error in this response without appealing to any such assumption, since the more fundamental problem with this version of the Impersonal View lies in the fact that perceptual experience is primarily a way of being aware of how things are *in the world*, whereas things like reference frames and points of view are rather features pertaining strictly to the structure of experience itself: these can be helpful concepts to draw on in certain kinds of psychological or phenomenological analysis, but such analyses regard them not as among the *contents* of experience but rather as aspects of their distinctive modes of presentation. This way of developing the Impersonal View forces us to suppose that visual experience includes among its contents not just the structure of the world but also the structure of experience itself: we must think that the point of view, or the origin point of the reference frame, is made manifest in visual experience *as* such a point, as opposed to a location occupied by a particular worldly thing, albeit one with a singularly distinctive role in the experiencing subject’s life. Perhaps this is a *coherent* thing to think, but it is hard to see why we should think it more palatable than simply supposing that visual experiences can be self-locating in a manner similar to the ordinary first-personal beliefs we form on their basis. Critical reflection on the Impersonal View thus leaves us with even more reason than before to believe that the Self-Location Thesis is true.

⁸ Another variant might propose that it is not “I”, but rather “the subject of experience” whose location is experienced in these cases, but the objections that follow would apply with equal or greater force to such a suggestion. Tom Avery, who argues for a version of the Impersonal View in his 2009, has proposed in conversation that the relevant content might be that “*here* is moving (/at rest)”, but of course it is precisely the content of the spatial indexical “*here*” that changes in a case of self-motion in a way that makes a sentence like “*Here* is moving” simply ungrammatical.

5 A Maximal View?

Our arguments against the Minimal and Impersonal Views have entitled us to the claim that *some* visual experiences have self-locating elements among their face value contents, but might it be possible for this argument to be extended to give us a much more general version of this conclusion? According to what we can call a *Maximal View* of self-locating visual experiential content, there are self-locating elements among the face value perspectival contents of *all* visual experiences (and perhaps even non-visual ones – though we will not explore that possibility here), rather than just the select range of cases relied on so far to argue for the Self-Location Thesis. In contrast to the Maximal View stands what we can call a *Modest View* of visual self-awareness, according to which the Self-Location Thesis holds true only in a restricted range of instances, but not universally. Having shown that the Self-Location Thesis is true of human visual experience at least some of the time, it remains to consider what the Maximal and Modest Views might have to be said in their favor.

To start, we should note that there are two different, logically independent dimensions along which a generalized version of the Self-Location Thesis might be thought to hold. First, it could be argued that just as we have argued for the presence of self-locating contents in experiences like E_1 , so it can be shown that *any* experience of the sort that ordinary human perceivers enjoy has perspectival spatial contents that are explicitly self-locating: the idea would be that *whenever* we describe the contents of an experience in terms that don't explicitly mention the perceiver's location there will be something we have left out; we simply do not have experiences that do not include our own locations among their face value contents. On the other hand, one might wish to extend the Self-Location Thesis beyond the confines of ordinary *human* experience and make it into a claim about the essential nature of visual experience in general: here the thought is that *whenever* we have a creature who is a subject of visual experience, some – or, if these two versions of the Maximal View are combined, perhaps *all* – of its visual experiences will have perspectival spatial contents that are explicitly self-locating.⁹ Clearly neither of these more ambitious theses entails the other, though as noted just above they can be combined into a view that is maximal along both dimensions. What do these theses have to be said for and against them?

In considering the second version of the Maximal View, according to which there can be no such thing as a visual experience whose subject cannot have at least some other visual experiences with explicitly self-locating content, the natural point to press concerns the operative conception of “experience”. For if, as was argued in Section 3.1, it makes sense to think of a system of spatial representation that represents only perceiver-relative, or perhaps both perceiver-relative and also “objective”, spatial position but has no resources to represent the location of the subject, then we need to get a better handle on why an organism whose visual perceptual contents were Modest in this sort of way would essentially fail to be a subject of visual *experience*. That is to say: granting that (1) *we* have visual experiences that are explicitly self-locating, and that (2) it is possible for there to be perceptual contents that are egocentric but do *not* include the location of the self, we need to understand what it is about the concept of “experience” that keeps it from applying to perceptual states with contents that are exclusively of the perceiver-relative but non-self-locating variety. It may be that a philosophically viable

⁹ Alva Noë, in conversation, has expressed some sympathy for this sort of position, and much of what I say about it below is indebted to our discussions.

position could be worked out along some such lines, but any of its details lie outside the scope of the present inquiry.

Thus for now we can restrict our attention to the version of the Maximal View according to which there are explicitly self-locating elements among the face value contents of all of *our* visual perceptual experiences. On this way of thinking, mature and psychologically normal human perceivers are not visually aware of their locations only *some* of the time; instead, it is *never* possible fully to characterize the contents of our visual experiences except in explicitly self-locating terms. Descriptions of the contents of visual experience in ways that don't mention the perceiver's location are never phenomenologically adequate; normal human visual experience is never non-self-locating in the way that a proponent of the Modest View would have it.

The chief philosophical difficulty facing this version of the Maximal View arises from what phenomenologists have termed the "experiential transparency" of the perceiving subject's body during attentive absorption in the worldly affairs that ordinarily occupy our thought and action. Sartre provides a compelling description of this phenomenon in *Being and Nothingness*:

I do not apprehend *my* hand in the act of writing but only the pen which is writing; this means that I use my pen in order to form letters but not *my hand* in order to hold the pen. I am not in relation to my hand in the same utilizing attitude as I am toward the pen; I *am* my hand. That is, my hand is the arresting of references and their ultimate end. The hand is only the utilization of the pen. In this sense the hand is at once the unknowable and non-utilizable term which the last instrument of the series indicates ("book to be read—characters to be formed on paper—pen") and at the same time the orientation of the entire series (the printed book itself refers back to the hand). But I can apprehend it—at least in so far as it is acting—only as the perpetual, evanescent reference of the whole series. Thus in a duel with swords or with quarter-staffs, it is the quarter-staff which I watch with my eyes and which I handle. In the act of writing it is the point of the pen which I look at in synthetic combination with the line or the square marked on the sheet of paper. But my hand has vanished; it is lost in the complex system of instrumentality in order that this system may exist. It is simply the meaning and the orientation of the system. (1992: 426)

On Sartre's account most of our ordinary perceptual experience is subordinated in just this way to our engagement with the things around us in terms of which our intentional goals are ultimately structured, and because this engagement has such an outwardly-directed focus our bodies tend simply to "vanish" from our conscious experience, as elements of the worldly "system of instrumentality" exhaust the contents of consciousness. Sartrean phenomenology seems to give us a reason to favor the Modest View over the Maximal one, by supplying instances of human perceptual experiences in which bodily self-awareness is absent.

But so long as we allow the possibility that there may be elements of sensory awareness that outstrip the objects of focal attention, we can reject Sartre's conclusion even as we acknowledge the general importance of the phenomenon he is describing: thus Shaun Gallagher and Dan Zahavi defend the claim of a general "tacit and non-thematic self-consciousness" by noting that many phenomenologists "reject the suggestion that we are *attentively* conscious of everything we experience" (2008: 55): thus for example a driver in traffic might be paying close attention to the way that the car in front of him is weaving out of traffic, while at the same time enjoying a "pre-reflective self-consciousness" of what his own experience is like; the subject is

aware that he is watching the car, but not “in the manner of paying attention to the watching” (ibid.). And the evidence for this, they suggest, is that when asked what he has “been doing, or thinking, or seeing, or feeling” at a given moment a human subject is “usually able to respond immediately, i.e. without inference or observation” (2008: 54), thus suggesting that he already had an awareness of the matters in question. Clearly the notion of a pre-reflectively self-conscious awareness of *one’s awareness*, which is Gallagher and Zahavi’s immediate concern, is different from that of the visual awareness of one’s spatial position – but is it possible for the Maximal View to be defended in a similar sort of way?

According to Joseph Schear the problem with Gallagher and Zahavi’s position is that the best argument they can give for it falls prey to a familiar objection:

We start with a point about what’s immediately knowable, hence reportable, if asked. The query (“what are you doing?”) is, metaphorically, the opening of the refrigerator. But just as it doesn’t follow from the light being on when we open the refrigerator that the light is always on, so it doesn’t follow from our being able to report knowingly on our conscious lives when asked that our conscious lives always includes self-consciousness. The fallacy is particularly inviting when we engage in phenomenological reflection. After all, to reflect on the structure and character of our *own* experience is an intensely self-conscious enterprise. As soon as we’ve set off on the investigation, we’ve “opened the refrigerator.” Unsurprisingly, self-consciousness turns up wherever we look. (Schear 2009: 101)

The idea here is that the mere ability to “respond immediately” at any given moment to questions about a certain subject matter is no good reason to think that that subject matter is something one must already have had an awareness of; instead, Schear proposes, it might just be that asking the “interview question” manages to “bring on a transformation of conscious experience—a shift of mental posture—rather than merely trigger the revelation of what had already been at work” (ibid.). Thus on Schear’s view we should think of self-consciousness not as a permanent aspect of conscious awareness but rather as “a potentiality—generally unactualized, but always actualizable—of the world-immersed experience of someone capable of first-personal thought” (2009: 99).

But can we really give a similar analysis of the phenomenology of visual self-location? One problem with using a variant of Schear’s argument in objecting to the Maximal View is that there is every reason to think that the kinds of visual phenomena we appealed to in arguing for the Self-Location Thesis can be present in cases of Sartrean absorption, i.e. cases where the perceiver’s attention is entirely captured by matters that have little or nothing to do with himself or his own location. For example, in Gallagher and Zahavi’s example of the driver who is caught up in watching the antics of the car in front of him there is no reason to think that the driver will be *entirely* unaware of his own motion, such that his visual experience is exactly the same as it would be if he were viewing the same scene but with his car sitting still: thus on their own account, despite the fact that we “do not normally monitor our [bodily] movements in an explicitly conscious manner” when we are absorbed in outwardly-directed behavior, nevertheless “we have a pre-reflective awareness of our body in very general terms”, such that even though “our attention, our intentional focus, is normally on the task to be performed, the project to be accomplished, or on some worldly event that seems relevant to our action”, the movements and position of one’s own body “are immediately and pre-reflectively felt” (2008: 145). And while Gallagher and Zahavi’s focus here is more on the *non*-visual proprioceptive and kinesthetic awareness of one’s body, what they say captures precisely the kind of role that a defender of the

Maximal View might envision for self-locating visual content: one's eyes, and with them one's mind, are in general directed outwards, but the experience they enable essentially carries with it a tacit awareness of, which is more than just a way of coming to know about, how things are with oneself.

There is, however, a further consideration that speaks in favor of the Maximal View even more forcefully than this first one, namely that rejecting the Maximal View on Sartrean grounds requires us to make sense of the extremely unintuitive idea that having one's attention drawn to the question of one's own location can bring about a "transformation of conscious experience" in which a continuous experiential episode alters from having merely perceiver-relative spatial contents at one moment to having explicitly self-locating ones at another. On this way of thinking about what happens when a perceiver "opens the refrigerator" and visually attends to his location, before the "interview" begins aspects of the perceiver's location are nowhere to be found in his experience, while after the question is asked and the perceiver shifts his attention appropriately the contents of his experience undergo a significant change, and self-locating aspects suddenly pop up: prior to the shift in attention the perceiver was aware of things only as being "to the right", "to the left", and so on; while after the "change in mental posture" the perspectival spatial contents of his experience change significantly, and he is now aware of where those things are with respect *to him*. And this is an odd proposal indeed: it is simple enough to spell out the idea of a visual experience that is perspectivally spatial but not explicitly self-locating, but much harder to make sense of a visual gestalt shift from an experience of this sort to one where the perceiver's location has suddenly shown up. At the very least, this consideration suggests that it is the Maximal View, and not the Modest one, that deserves to be our default position on the issue of visual self-location; the burden of proof is on the proponent of the Modest View to provide positive phenomenological evidence *for* the kind of transformation envisioned here, rather than the other way around.

But how exactly does adhering to the Maximal View require us to think about the nature of visual self-awareness? Part of what Gallagher and Zahavi are out to achieve with their talk of "pre-reflective" self-consciousness is to differentiate the "objectifying" perceptual experience we have of the objects around us from the first-personal awareness of oneself "as a subject", as what Husserl calls a *Leib* in contrast to a merely physical *Körper*; and as we saw in Section 1 of the Introduction one natural way to effect this distinction is to hold that self-awareness is primarily constituted through some sort of "inner sense". Does the claim that we are aware of ourselves "from the outside", as it were, require us to abandon the idea that we have a way of knowing about ourselves that others do not? Can we make sense of the possibility of a visual self-awareness that is not a queer sort of self-*observation*? The aim of my concluding chapter is to argue that we can.

Chapter Three: Perception and Practical Knowledge*

1 Introduction

Philosophical thinking about intentional action can tend to pull us toward two deeply intuitive but fundamentally incompatible positions. On the one hand, if we begin from the idea that the knowledge a person has of his own intentional actions is essentially different both from his knowledge of the actions of others and of any actions of his own that are involuntary or unintentional, then we can quickly be led to think that the operation of intentional agency is fundamentally an “interior” affair, having at its core some domain of events that a person is guaranteed to have a way of knowing about when they are a part of his own life that he cannot draw on when they are a part of the lives of others. If, however, the action theorist takes his departure from the idea that, as Elizabeth Anscombe puts it, in acting intentionally “I do what happens” (2000: 52), i.e. that human actions are events every bit as “worldly” as the rolling of a stone and the fall of a sparrow, then a natural place to end up is with the conclusion that one’s knowledge of his own intentional actions can’t be that different from his knowledge of other sorts of events after all, or at least that the epistemic specialness indicated earlier isn’t characteristic of the entirety of one’s intentional behaviors. Clearly these conclusions can’t both be right, yet each of the starting points highlights an aspect of intentional agency that seems entirely beyond dispute: is it possible, then, to occupy a position that acknowledges the force of those initial claims but rejects the corresponding conclusions?

These tensions in our philosophical understanding of agency neatly parallel the divergent demands for a theory of self-awareness that were briefly canvassed at the end of Chapter Two. For thinking of the spatial content of visual experience as essentially self-locating along the lines proposed by the Maximal View can lead naturally to the idea that human self-awareness isn’t experientially distinctive in any philosophically important way; while starting from the idea that there *must* be something *sui generis* in the structure of self-awareness leads to the idea that visual self-location doesn’t undergird a *true* form of self-awareness after all, at least not one on a par with the introspective awareness of one’s mental states and the awareness of bodily position through dedicated sensory channels.

It is helpful to think about Anscombe’s *Intention* as, at least in part, an attempt to show the way out of these sorts of dilemmas. For in the first place Anscombe is unwavering in her insistence that a person’s knowledge of his own intentional actions is a distinctive kind of knowledge, both in being “non-observational” and in being a species of “practical knowledge”, knowledge that is somehow “the cause of what it understands”. Yet at the same time, she flatly rejects any attempt to reduce intentional agency to, or indeed even to treat it in a way that privileges our claim to have a special knowledge of, such (putatively) “interior” events like intendings, tryings, or “mere” bodily movements. If Anscombe is right, then we can respond to the dilemmas sketched above by grasping both horns in each case; we can preserve what is distinctive about human self-awareness without reducing its subject matter to a restricted range of inner events.

This chapter aims to show how the Anscombean idea of non-observational knowledge can be put to work in explaining how the knowledge of one’s own intentional actions differs both

* Thanks especially to Randall Amano, John Campbell, Ben Kiesewetter, Niko Kolodny, and Alva Noë for helpful conversations regarding various aspects of this material.

from the knowledge that others can have of those actions and from the knowledge a person has of what he unintentionally does. It also attempts, however, to show that several other recent attempts to work out this sort of position fall short of their goal by mischaracterizing certain key features of intentional actions and the kind of knowledge agents have of them. I begin in Section 2 by sharpening the problem, considering some recent attempts to resolve it, and arguing that none are fully successful. Section 3 develops and provides support for an alternative response, and in Section 4 I consider the bearing on the philosophical position developed here of empirical research on the relationship between conscious vision and the online guidance of action. It is, I will argue, not because it has a special source but rather because of the distinctive role that it plays in our mental and behavioral economy that the knowledge we have of our own intentional actions is epistemically distinctive, and that it deserves to be called non-observational even as sense perception plays a fundamental role in making it possible.

2 The Self-Knowledge Problem

How should we think about the way in which a person's knowledge of his own intentional actions is fundamentally different from whatever knowledge he might have of other events, especially his unintentional behaviors and the intentional and unintentional actions of others? The common answer given by philosophers staking out an Anscombean position is that it is because only knowledge of the former sort is essentially *non-observational*, and a standard assumption in the literature on this topic is that the notion of non-observational knowledge is best understood as picking out a kind of knowledge that is grounded independently of sense perception. Thus Kieran Setiya writes that "when an agent is doing ϕ intentionally, he knows that he is ϕ -ing, and he knows this spontaneously, not on the basis of empirical evidence" (2008: 392); and in general, though this assumption is rarely made as explicit as Setiya makes it here, the idea that independence from observation requires epistemic independence from the deliverances of sense perception is shared by defenders and critics of Anscombe alike.

While I will argue in this paper that our knowledge of our intentional actions is *not* independent of observation in this sense of the term, and that the most promising way to develop this aspect of Anscombe's view is to understand the notion of non-observational knowledge in quite another way, it is impossible to deny that Anscombe, too, comes close to identifying "observational" knowledge with knowledge grounded in sense perception. When the notion of non-observational knowledge is introduced near the start of *Intention*, it is by appeal to the difference between bodily awareness and coming to know about things through, as she puts it a bit later on, the "exterior senses" (2000: 49); and when she tries to bring out the specialness of agents' knowledge of their own intentional actions it is usually with examples like the following:

Say I go over to the window and open it. Someone who hears me moving calls out: What are you doing making that noise? I reply 'Opening the window'. I have called such a statement knowledge all along; and precisely because in such a case what I say is true—I do open the window; and that means that the window is getting opened by the movements of the body out of whose mouth those words come. But I don't say the words like this: 'Let me see, what is this body bringing about? Ah yes! the opening of the window'. Or even like this: 'Let me see, what are my movements bringing about? The opening of the window'. To see this, if it is not already plain, contrast this case with the following one: I open the window and it focuses a spot of light on the wall. Someone who

cannot see me but can see the wall says ‘What are you doing making that light come on the wall?’ and I say ‘Ah yes, it’s opening the window that does it’, or ‘That always happens when one opens the window at midday if the sun is shining.’ (2000: 51)

As we will see shortly, the standard way of interpreting these sorts of examples is to say that it is because the agent has to *perceive* that he is casting the light on the wall in order to know that he is doing it that casting the light cannot be counted among his intentional actions, while he doesn’t rely on sense perception in this way to know that he is opening the window. It is for this reason that, according to Anscombe, replying “I knew I was doing that, but only because I observed it” to a request for a reason-giving account of one’s action, as for example “if one noticed that one operated the traffic lights in crossing a road”, is a way of implying that the action in question was not intended (2000: 14): having to “observe” or “notice” that you are doing something in order to know that you are doing it rules out the possibility that it is among the things you intentionally do.

Nevertheless Anscombe is acutely aware of the intuitive difficulty in understanding how things like my opening a window can be among the things I know in such a special sort of way, and she anticipates the following objection:

‘Known without observation’ may very well be a justifiable formula for knowledge of the position and movements of one’s limbs, but you have spoken of all intentional action as falling under this concept. Now it may be e.g. that one paints a wall yellow, meaning to do so. But is it reasonable to say that one ‘knows without observation’ that one is painting a wall yellow? And similarly for all sorts of actions: any actions that is, that are described under any aspect beyond that of bodily movement. (2000: 50)

Clearly there are lots of different behaviors that suggest concerns along these lines: painting a wall yellow is one; opening a window is another; and so are casting a light on a wall, reading a magazine, lifting a box, chopping an onion, putting ham on a sandwich, and so on. In each case the intuitive difficulty with treating your knowledge of your actions as non-empirical arises from the fact that (1) it is impossible actually to engage (and not merely: try to engage) in an action of the type in question without effecting changes in the world that involve more than the movements of one’s body, but (2) an agent will sometimes be able to know only by sense perception that he is bringing such extra-bodily changes about, and so (3) it appears that an agent will sometimes know only by sense perception that he is acting as he intends. Anyone who wants to use independence from sense perception as the criterion defining the privileged access we have to our intentional actions thus needs to reject one or both of premises (1) and (2), or explain how the inference from these premises to the conclusion in (3) is invalid.

Call the problem of explaining how, in cases like these, an agent can know without observation what he is doing under all of the descriptions under which he is acting intentionally the *Self-Knowledge Problem* for Anscombe’s theory of intention. Later on in this chapter I will attempt to articulate a conception of non-observational knowledge that I think can resist the objection, but first I want to consider a few recent responses to it and show why each of them fails.

2.1 Rejecting (1)?

According to the first premise of our argument for the Self-Knowledge Problem, it is sometimes impossible to act in a given way without bringing about changes in the world at a distance from one's body: it is in this sense that "what one does" *just is* "what happens"; and consequently it seems impossible to know whether one is acting in such a way without knowing whether such extra-bodily happenings really are taking place. Yet there are many cases in which our commonsense metaphysics of action allows us to think of people as engaged in certain actions even when they are *not* bringing about the sorts of worldly changes that would be required for those actions to be completed: for example, I can spend many years working on a dissertation even as nothing that deserves that title manages to materialize; and moreover I can be said to be working on it at those times when I am not presently engaged in any real "work" at all, and even when the status of my dissertation project could not be further from my mind. Thus Anscombe:

A man can *be doing* something which he nevertheless does not *do*, if it is some process or enterprise which takes time to complete and of which therefore, if it is not cut short at any time, we may say that he *was doing* it, but *did not do* it. This point, however, is in no way peculiar to intentional action; for we can say that something was falling over but did not fall (because something stopped it). Therefore we do not appeal to the presence of intention to justify the statement 'He is Y-ing'; though in some cases his own statement that he is Y-ing may, at a certain stage of the proceedings, be needed for anybody else to be able to say he is Y-ing, since not enough has gone on for that to be evident; as when we see a man doing things with an array of wires and plugs and so on. (2000: 39)

Noting this feature of our action concepts can incline us to think that whether or not one acts in a given way doesn't actually require effecting the kinds of worldly changes that gave rise to the Self-Knowledge Problem: if I can act in a given way *without* bringing about the extra-bodily events that constitute a completed act of the sort in question, then it seems that I don't need to have knowledge of such events in order to know what I do.¹ This view is compatible with the idea that one very often has to accomplish *something* extra-mental in order to act, but perhaps the idea is it is possible to engage in a certain kind of action, say one of ϕ -ing, not just by bringing about the happenings that a *completed* act of ϕ -ing involves but also by doing something else, say ψ -ing, that is a *means* to such a completed act, where one's ψ -ing can consist of happenings that take place at less of a distance from one's body than would those that constitute the fully completed act of ϕ -ing. And since (one might suppose) it is possible to have non-observational knowledge that one is acting in the "proximal" manner, we can know in this way that we are doing all of the various things we are intentionally up to.

But this is not Anscombe's view at all, and in fact it runs dangerously close to what she calls the "false avenue of escape" according to which "I really 'do' in the intentional sense whatever I think I am doing" (2000: 52). According to Anscombe the problem with this way of thinking is that it is impossible to say how the "vehicle" for an intention could be independent in this way of a worldly occurrence that is the achieving of what one intends; but whether or not we are persuaded by her response this objection to premise (1) is prey to a more immediate problem, namely that our ordinary intuitions about what it takes to engage in an action are not nearly as liberal as the objection supposes. As Kevin Falvey notes:

¹ For an example of such a view, see Paul MS: 23-26.

The precise details of the truth-conditions of sentences in the progressive [tense] is a matter of considerable obscurity, but it is surely a necessary condition for me to be φ -ing that it be possible for me to φ . Hence the claim that I am or was φ -ing must be withdrawn if in the course of things it becomes clear that I cannot φ ; either because I lack the general skills or ability to φ , or because I lack the materials required for φ -ing. If I put the kettle on the stove and say, “I’m making tea,” and am told that the stove isn’t working, I should not subsequently say that I was *making tea* when I learned that the stove wasn’t working, I would retreat to something like, “I was going to make tea” (or, “I was trying to make tea”). On the other hand, if, as I am putting the kettle on for tea, the phone rings and the tea-making is aborted, I could still later say that I was making tea when my friend called, as a result of which I didn’t make tea. The openness of the progressive allows for interruptions of actions-in-progress, including changes of mind. But from the fact that an event or process of a given type could not have been completed in the circumstances, it does seem to follow that no event of that type could have been underway. (2000: 24)

Falvey is right, and in fact the point applies more widely than those instances where the sorts of actions under consideration are simply impossible: to return to our original example, it seems that a person who means to paint a wall yellow but is accidentally using white paint instead simply doesn’t count as painting the wall yellow on our ordinary understanding of what that action involves; he’s *trying* to paint the wall yellow, we might say, and while he’s doing some of the things that are required actually to do that it’s nevertheless not the action he’s really engaged in. So premise (1) seems well enough in order: there are some cases where acting intentionally in a certain way requires bringing about specific changes at a distance from one’s own body, and the crucial question concerns whether an agent can know without observation that those changes are being brought about.

2.2 Rejecting (2)?

If human agents essentially possess non-observational knowledge of their intentional actions even under descriptions involving changes that happen at some distance from their own bodies, and if what it is for something to be non-observationally known is for it to be known independently of sense perception, then it seems that whenever someone is intentionally engaged in an action that requires bringing about certain changes in the world he must independently of sense perception that he is bringing such changes about. Attempting to explain how such knowledge is possible is the most common way that defenders of a broadly Anscombean theory of action have responded to the Self-Knowledge Problem, and this section considers a couple such attempts and argues that they fail to show where our premise (2) goes wrong.²

2.2.1 Knowledge by “Know-How”

Kieran Setiya’s “Practical Knowledge” argues that the Self-Knowledge Problem vanishes once we recognize two things: first, that the intention to act in a certain way consists in part in the *belief* that one will do just that; and second, that when we intend to do things we generally *know*

² For a further objection to accounts like those that I discuss in the following sections see Gibbons MS: 14-17.

how to do them, and moreover know that we have such abilities. According to Setiya, once these factors are in place there is no trouble in seeing how exercises of human agency essentially involve non-empirical knowledge of what one intentionally does: for assuming that an action is intentional if and only if the agent intends to engage in it, then the first principle yields the claim that the agent will also *believe* that he is engaging in it, while the second principle yields the further conclusion that this belief will be justified by his knowledge that he'll be able to do it if he so intends. As he puts it:

Knowing how to ϕ is the state or condition that, with knowledge of ability, provides the epistemic warrant for decision. Together, they justify the transition in which one forms the intention and belief that one is doing ϕ or that one is going to do it. More carefully, this transition is justified if and only if one's decision is an exercise of knowledge how to ϕ and one has knowledge of ability, in the simple conditional sense. Knowledge how thus plays a role in dynamic epistemology, in our entitlement to form and revise beliefs. (2008: 407)

By "the simple conditional sense" in which we have knowledge of our abilities Setiya means the sense in which a person can know that, if he intends to be doing something at some time, then at that time he will be doing so in fact. If we possess knowledge of this sort whenever we act intentionally, and if such actions are always accompanied by justified beliefs in their occurrence, then will it follow that our knowledge of what we intentionally do is essentially independent of sense perception?

The fundamental problem with this position is that it leaves untouched the possibility that even a perfectly able agent will sometimes *fail* to do what he intends, and when the possibility of such failure is sufficiently real we are left without a way to understand how the agent knows that such a possibility does not obtain. For even Anscombe, despite insisting that "the failure to execute intentions is necessarily the rare exception" (2000: 87), allows that intentions *sometimes* "fail to get executed" (2000: 82); and clearly it is possible to act intentionally in a given manner even in instances where, one's general ability so to act notwithstanding, the surrounding circumstances render it a real epistemic possibility that one's intended action might fail to materialize. If, for example, I am presently painting a wall a perfectly uniform shade of yellow, then even though I may (i) believe that I am painting it in such a way and (ii) know that I am *able* so to paint it, nevertheless I may not (iii) *know that that is the way I am painting it* without seeing how the paint is actually going on. Having the ability to ϕ simply cannot require being able to ϕ successfully whenever I see fit to try; and even if it did then Setiya's account would fail to apply to the vast majority of our intentional actions, since of course we don't generally take ourselves to be able to act on our intentions with nearly this high a rate of success. It may be that intentional agency always requires knowledge of ability, but that alone is not sufficient always to justify us in believing that we have followed through on our intentions, and so cannot provide the full account of our first-personal knowledge of our actions.

2.2.2 Inference from Intention

According to Sarah Paul's neo-Gricean "inferential" account of how we know our own actions, the knowledge of what one is intentionally doing is based not on sense perception but rather on an (often unconscious) *inference* from the knowledge of what one intends, plus empirical and

non-empirical knowledge of the background conditions that make success likely: “what the agent knows evidentially”, she writes, “is what he intends to be doing, while insofar as he has a belief about what he is actually doing, this is evidentially based on his knowledge of what he intends, plus his evidence for thinking that he will do what he intends” (MS: 17). Thus unlike Setiya, Paul rejects the idea that intending to act in a certain way *essentially* involves the belief that that is how one will act; she retains the commitment, however, to the idea that the justification for an agent’s belief that he is acting intentionally in a given way is independent of any perceptual awareness of that action.

But Paul is clearly sensitive to the sort of complaint we raised above against Setiya’s view, that merely knowing that I *can*, even most of the time, act in a certain way if I intend to will not always justify me in believing that I *am* so acting on any given occasion. Thus she allows that the agent’s “awareness of his circumstances as being conducive to his ϕ -ing (or at least, not being likely to obstruct his ϕ -ing)” will be a part of what grounds his knowledge of what he does (MS: 22): such awareness may be observational, of course, but Paul claims that since it “is not experience or observation of the particular action in question” (MS: 23; emphasis added) it will not be the sort of perception-based knowledge that Anscombe’s account is meant to rule out. And so according to Paul it is because *this* sort of knowledge usually accompanies the knowledge of one’s intentions and abilities that the knowledge of one’s own intentional actions is non-observational: given that I know (i) what I presently intend to do, (ii) that I usually do what I presently intend, and (iii) that I am capable of doing just that in circumstances like these, I am justified in believing that I am presently acting just as I intend to be.

How should we think about the viability of Paul’s inferential view as a response to the Self-Knowledge Problem? In the first place, it seems open to question whether this account of non-observational knowledge really manages to preserve the first-personal/third-personal asymmetry that was emphasized at the beginning of this chapter, i.e. the way in which my knowledge of my own intentional actions seems essentially different from your knowledge of them. For on an account like Paul’s while there will clearly be fundamental differences between our respective ways of knowing about my *intentions*, once that knowledge is in place it seems as if you will be in every bit as good a position as I am to infer what I am doing from those intentions plus the knowledge of my general tendencies and the favorability of my circumstances: this is an inference that you are every bit as capable of making as I am, and it seems that you will be no less justified. Yet it seems as if it is knowledge of one’s *actions*, and not just of the intentions that give rise to them, that ought to be characterized by first-personal privilege according to Anscombe’s very intuitive picture, and it is a significant defect of Paul’s position if it fails to ensure this.

Moreover, on closer inspection Paul’s account does not fare much better than Setiya’s when we consider the question of how we know what we are doing in sufficiently *unfavorable* circumstances. If, for example, I am successfully steering my boat through a narrow passage, my perceptual awareness of the storminess of the sea may give me reason to believe that the surrounding conditions are not especially conducive to steering the boat as I intend to, and so to know that I am steering toward the intended passage rather than (unintentionally) heading for a dangerous pile of rocks I will need to rely on my perceptual experience of *where I am going*: yet if asked why I am steering the ship through the passage rather than toward the rocks it would be stupid of me to reply that “I knew I was doing that, but only because I observed it” (Anscombe 2000: 14). If relying on sense perception in this sort of way to know that I am acting, and not merely trying to act, as I intend is incompatible with my having non-observational knowledge of

what I am doing, then Anscombe's claim that we *always* know our intentional actions in a non-observational manner is clearly false. It may be that Paul's account is sufficient to explain how we can be justified in believing that we are acting as we intend some or even most of the time, but it cannot apply as generally as Anscombe's account is meant to.

2.3 Rejecting the Inference to (3)?

Once we accept premises (1) and (2) in our argument for the Self-Knowledge Problem it might seem that the conclusion in (3) follows immediately; yet Richard Moran has staked out a position that can be viewed as an attempt to deny just this. According to Moran, what allows an agent's knowledge of his own intentional actions to "involve reference to actual changes in the world" beyond the confines of his body (as in premise (1)) but still differ from an outside observer's "speculative, observational knowledge of what is the case", is that the agent's knowledge "commits itself not only to the obtaining of certain events in the world, but to the specification of the descriptions under which what happens counts as the execution of his intention" (2004: 56). The guiding idea here is that our understanding of intentional action is essentially an *intensional* affair: "to say", as Anscombe puts it, "that a man knows he is doing X is to give a description of what he is doing *under which* he knows it" (2000: 12). And Moran's proposal is that the difference between the non-observational knowledge of one's own actions and the observational knowledge of merely physical happenings or of the actions of others is that an agent's knowledge of the descriptions under which his actions are truly characterized as intentional is essentially of the former sort:

Understood extensionally, I can know *what happens* only through observation, including the perception that serves as an aid in the execution of action such as writing on the blackboard ..., and the causal knowledge Anscombe refers to earlier as "knowledge or opinion concerning what is the case, and what can happen—say, Z—if one does certain things, say ABC" ... With this empirical knowledge in place, I can form an intention to *do* something, such as opening a window, and then actually *do that thing*. But the event which is my action only counts as something I intentionally do in virtue of some of its descriptions and not others, and my knowledge of it can be said to be 'non-observational' *only* under the terms of such descriptions. (2004: 55-56)

So unlike Paul and Setiya, Moran allows that my knowledge that I am actually effecting the sorts of changes in the world that are required for me to count as acting as I intend will sometimes involve a sense-perceptual awareness of what I am bringing about. Yet he wants to resist the idea that this means my knowledge *of my actions* can't meaningfully be classified as non-observational, since my "empirical knowledge" of what happens in the world won't be enough to settle the descriptions under which those happenings count as an intentional action, as this latter sort of knowledge "depends on, but does not reduce to, the speculative knowledge of what can happen and what is happening, and in this way Anscombe may evade the charge ... that the admitted dependence of successful action on ordinary observation must mean that the agent's knowledge of what he is doing is ultimately observational after all" (Moran 2004: 56).

But is it really true that "observational knowledge of what is the case" essentially fails to specify the descriptions under which the observed happenings count as an intentional action? Moran motivates this idea by appealing to an example of David Velleman's (1989: 15), in which

a person walking through Manhattan realizes that he doesn't know what he is doing. Of this example, Moran writes:

... such cases enable us to see how it can be true both that the agent will normally depend on observation of various kinds to carry through an action like walking up Fifth Avenue ..., and also that the knowledge that is temporarily lost in such cases is nonetheless not made up for my further observations alone. The person stops and *looks around*, observing his position and his environment for clues as to what he might have been up to, but this by itself does not deliver to him the knowledge of what he is or was doing, for it does not provide him with the particular set of descriptions of what he sees or the movements he is making, under which what he was doing counted as an intentional action of his. (2004: 57)

This example is supposed to help us see, Moran continues, “how practical knowledge *could not* be observational, could not be perceptually derived from the world. For nothing the agent sees in the world could give him those descriptions, even though *what* is claimed in practical knowledge is a world-involving matter of fact” (ibid.). Yet this conclusion – that the knowledge of what one is intentionally doing requires a component which observation cannot possibly supply – only follows from this example if the walker’s perceptual experience *is the same* in both cases, i.e. when he is walking absent-mindedly and without knowledge of what he is doing and when he is walking intentionally and with full self-knowledge: and it is only if we make the implausible assumption that intentional and absent-minded behaviors alike are experienced through sense perception only as sets of bodily movements that we would be entitled to treat these two cases in such a way. Moran is certainly right to say that “further observation of the world” (2004: 58) will not provide the man who stops and asks what he is doing with the right sort of answer to the question of what he is doing, but it does not follow from this that “the knowledge he has temporarily lost is something that goes beyond his observational knowledge” (ibid.), i.e. that when he *did* know what he was doing, this knowledge had a non-empirical component: for this would be a valid inference only if what was in the world to be observed was the same when he was acting intentionally as when he was not, but clearly there is no good reason to think that this is so.³ Otherwise, we are perfectly entitled to say that the reason sense experience does not provide the absentminded walker with knowledge of his intentional actions is that in his case there are no such actions to be experienced at all, while for the person presently exercising intentional agency, things are experientially quite different, as his intentional behaviors are right out in the open.

Moreover, even if this objection were to be rebutted Moran’s account is one on which one’s knowledge of his intentional actions is non-observational only in a rather weak sense, one which runs perilously close to a “two-factor” view according to which the only thing *really*

³ As Moran himself puts it: “There is an ordinary sense in which [the absent-minded walker] *sees* the same things as he did when he was engaged in action (e.g., walking up Fifth Avenue). But there is also a sense in which what he sees is now blank to him, because he cannot articulate his own relationship to it. The objects and scenes of his environment no longer have a role assigned to them in his ongoing action (as goal, obstacle, distraction, background, etc.). He sees Fifth Avenue, and he sees it from a particular perspective which indicates what direction he was heading in, but he does not see his goal. Instead, he sees *everything* around him, the traffic going by, the meaningless street signs, the strangers’ faces, and nothing in those details enables him to discern a destination, a point to his being right here facing in this direction” (2004: 57). Some similar points are made by Brian O’Shaughnessy (1980: 19-21).

known by the agent without observation is something like what he intends.⁴ It seems that what Anscombe is really after is a view according to which *even what happens* in a case of intentional action is something that an agent knows about in a special sort of way: it is not just that I will always know in a way that you cannot the descriptions under which my bodily behaviors count as intentional actions, but also that I will have a special epistemic relationship to those happenings even in the “extensional” sense, even as they involve “mere” changes in the physical world. It is this that Anscombe seems to be after when she writes that “in so far as one is observing, inferring etc. that [some event] Z is actually taking place, one’s knowledge [of Z] is not the knowledge that a man has of his intentional actions” (2000: 50); and it is hard to see how Moran’s analysis allows this diagnosis to come out true. In the next section of this chapter I will argue that it is by developing a conception of non-observational knowledge that requires something other than independence from sense experience that we can resolve the Self-Knowledge Problem, and do the most justice to the spirit of Anscombe’s view.

3 Practical Knowledge

According to the argument of Section 2 it is impossible to reconcile our commonsense metaphysics of human action with the claim that our knowledge of what we intentionally do is essentially independent of sense perception: there are, we saw, plenty of possible cases where the an agent’s perceptual awareness of his actions plays an important role in grounding his knowledge of what he is intentionally doing. But how can this conclusion be squared with the deeply intuitive idea that an agent’s knowledge of his intentional actions is epistemically distinctive, differing both from his knowledge of other sorts of things and from the knowledge of his actions that might be had by an outside observer?

Well, what exactly do we mean in this context when we speak of an *observer*, and of knowledge had *by observation*? Surely the implication of a reliance on sense perception is important here, as part of what it is to observe something is to become aware of its nature through perceptual experience. Yet the notion of observation carries further connotations than this: in particular, another aspect of observation is the perceiver’s relative *passivity* toward the objects or events in question; to observe something is to sit back, as it were, and simply take it in for whatever it happens to be. As Brian O’Shaughnessy puts it:

The conductor cannot listen as *observer* to the music he makes, since he already listens to it from the standpoint of *creator*. His listening is logically subordinated to that activity, to which it relates somewhat as the painter’s looking relates to the activity of painting. Were he to listen as observer, his listening would no longer be co-ordinated with and logically subordinated to the act of making music. In sum: one cannot be listening to the music one makes both from within the act and from without the act; one cannot simultaneously listen in two different ways. Now this is the nature of the difficulty where the putative observer-sense is playing an essential stage-setter role for the activity it putatively studies. (1980: 29)

Clearly O’Shaughnessy is not denying the obvious fact that in a case like the one he describes the conductor can and will *hear* the music his orchestra makes; indeed, it is only because he does

⁴ For one statement of such a view see Donnellan 1963.

hear it that he can go on conducting, which is what O'Shaughnessy means by speaking of the "essential stage-setter role" that the conductor's auditory perception plays in his ongoing activity. Rather, the idea here is that there are certain *ways* of listening or otherwise attentively perceiving that are impossible when the subject-matter of perception is something the perceiver himself is intentionally doing: in trying to take such a dual attitude toward what he does, O'Shaughnessy says, the (putative) agent-observer is trying to be both "within" and "without" the very same action (1980: 31-32), and this is a stance no person can possibly adopt.

Put somewhat differently, the point O'Shaughnessy is making in denying that a person can ever observe his intentional actions is not that human agents are never sense-perceptually aware of the things they intentionally do, but rather that the self-awareness characteristic of intentional agency is essentially bound up in an exercise of self-*control*. What makes the observational mode essentially a passive one is that, while it may embody a preference for things to turn out one way rather than another, the observer is not the one who is bringing those changes about; and thus the experienced world of the observer "is one that is going its own way, is taking its own course" (O'Shaughnessy 1980: 20). By contrast, the agent will keep track of his actions not only with a purpose in mind but also in such a way that it is more than an act of fate whether things turn out accordingly; the agent is a "creator", and his awareness of his actions is subordinated to the task of ensuring that things proceed along the course he intends.

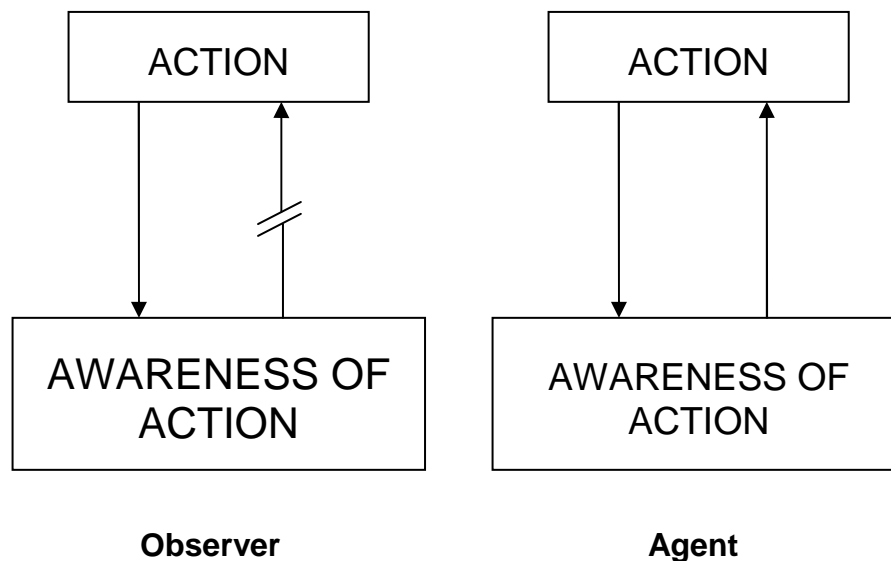


Figure 3: O'Shaughnessy's distinction between the agent and observer models of action awareness.

O'Shaughnessy's distinction between the kinds of action awareness characteristic of agency and observation respectively is illustrated very schematically in Figure 3. For our immediate purposes, the crucial point to note about it is that if the distinction between observational and non-observational knowledge breaks down along these lines rather than hinging on whether the knowledge in question is empirically grounded, then there is reason to hope that the Self-Knowledge Problem as articulated in Section 2 will be no problem for us at all; we will be able to reject the assumption that just because we sometimes have to rely on our

sense-perceptual awareness of our actions to know what we are intentionally doing, our knowledge of those actions therefore requires us to go in for self-*observation*. The remainder of this section will make the case for such an account of the non-observational knowledge of one's intentional actions, arguing that the best way to develop O'Shaughnessy's is by way of Anscombe's notion of "practical knowledge".

3.1 "The Cause of What it Understands"

According to O'Shaughnessy's account of the kind of action awareness characteristic of intentional agency, the distinguishing feature of such awareness is less a special source than a distinctive causal relationship that it has to its subject-matter. To emphasize this aspect of an agent's awareness of his actions is not, of course, to deny that that awareness *has* special sources: for a human agent is, at least ordinarily, aware of the movements of his body not just through ordinary sense perception but also through intra-bodily channels that give him a distinctive way of telling where his limbs are at any given moment. But Anscombe's account of intentional agency invites us to reject the idea that the involvement of such privileged sensory channels in an agent's awareness of his actions could constitute the *entirety* of what makes that awareness distinctive; and thus she gives us cases like the following:

... suppose someone simply wanted to produce the effect that in fact I lowered my arm at the speed at which it would fall—he is a physiologist, and wants to see if I generate anything different in my nerve fibres if I do this. So he fixes up a mechanism in which something in motion can be kept level if I hold a handle and execute a pumping movement with my arm and on the downward stroke lower it at the rate at which it would fall. No my instruction is: Keep it level, and with a bit of practice I learn to do so. My account of what I am doing is that I am keeping the thing level; I don't consider the movement of my arm at all. I am able to give a much more exact account of what I am doing at a distance than of what my arm is doing. So my keeping the thing level is not at all something which I calculate as the effect of what I really and immediately am doing, and therefore directly know in my 'knowledge of my own action'. In general, as Aristotle says, one does not deliberate about an acquired skill; the description of what one is doing, which one completely understands, is at a distance from the details of one's movements, which one does not consider at all. (2000: 54)

On Anscombe's analysis, when she keeps the thing level she has non-observational knowledge that this is what she does; yet clearly her knowing what she does cannot be explained just in terms of her reliance on bodily proprioception in what Shaun Gallagher calls "the ordinary (non-visual) sense of somatic (mechanical) information about joint position and limb extension" (2005: 46), for while there may be *some* aspects of how she is moving her body that Anscombe comes to know about in this sort of way, knowing that one is moving one's body is clearly not a way of knowing that one is keeping something level, and it is under the latter description that Anscombe's action is intentional. Nor, given the argument of Section 2, can we account for Anscombe's knowledge of what she does in terms that do not implicated sense perception at all: for the action of keeping the thing level is one that requires Anscombe to achieve a certain effect in the extra-bodily device, and if the task is sufficiently difficult then it may be that her knowledge of that effect is based on her perceptual awareness of what she does. But given

O'Shaughnessy's account of the observational/non-observational distinction this is no threat to the idea that Anscombe knows about her action without observing it: she is *aware* of her action, to be sure, but what makes this awareness distinctive is not its special source but rather the role it plays in shaping the unfolding of the very action it is an awareness of.⁵

It is not, however, only the agent's awareness of his action that has a special role to play in the way that action unfolds; on Anscombe's account there is also a way in which our *knowledge* of our own actions is causally relevant in bringing them about. In developing this conception of "practical knowledge" she references a passage from the *Summa Theologiae*:

... man is more like God with respect to his practical intellect, which is the cause of things thought of, than his speculative intellect, which derives knowledge from things. Therefore man's happiness consists in activity of his practical rather than his speculative intellect. (Aquinas 1983: 34)

Aquinas goes on to reject this analysis of human happiness as giving insufficient weight to the nobility of speculation, though he agrees that "the practical intellect is related to what it knows as God to what He knows" (1983: 35): human agents are knowers not just by taking things in but also in the ways they bring things about. When Anscombe keeps the device level by operating the pump handle, for example, we explain her actions partly in terms of her knowledge of how those very actions unfold: it is because she knows that she is keeping the thing level that she proceeds along as she has been; whereas if she comes to believe that the thing is off kilter she will reshape her action accordingly.

Yet quite *unlike* God, it is not simply by saying the word that a human agent can do what he wills, and thus it is generally only by virtue of the appropriate sort of sensory awareness that the knowledge of our actions will be fully in place. Not all of the feedback relevant to behavioral control proceeds by way of what we *know*, of course: this is one of the lessons of the example cited just above, where what the agent operating the pump handle "completely understands" is only whether the thing is being kept level, and in doing this she makes implicit use of proprioceptive feedback but takes little or no account of how she is moving her limbs. But as we noted, even in this case there is *something* which the agent knows about, namely the status of the thing she is trying to keep level, and it is partly because that empirical but practical knowledge is in place that her action unfolds in the particular way it does.

Thus according to the account on offer here our practical knowledge *itself* implicates a kind of knowledge that is "derived from things", even if it is not "speculative" in the ordinary sense of the notion: we are like God in being able intentionally to bring things about, yet thanks to our finitude we generally have to rely on sense perception both in acting and in having knowledge of whether or not we act as we intend. For it is only very rarely that a human agent intentionally engages in an action so simple that he can successfully achieve, much less know that he achieves, what he intends without paying any attention at all to what happens. In less

⁵ For some similar points see Roessler 2003, which distinguishes the agent's use of perceptual attention to answer "practical question" from the observer's use of it to answer "theoretical" ones. Yet on Roessler's account the answer to a practical question is *not at all* a theoretical matter; its content is not a way of staking a claim as to how things are in the world. A more natural view, I think, is that the agent relies on perceptual experience to take in how things are and immediately implicates this knowledge in deliberately shaping those things in the intended ways; there is no reason why the relevant deliverances of sense perception cannot have a mind-world direction of fit. But further discussion of this point would take us too far afield.

simple cases, the agent’s awareness of his behaviors and their effects on the world is part of a feedback loop that shapes and is in turn shaped by the ways his actions unfold. It is because the conductor can *hear* the music that the instrumentalists are producing that he knows how things presently stand and where he needs to proceed from here; and if his knowledge were other than it is – if, say, he misheard which notes the strings were playing and so lost track of where in the performance he was – then the course of his action would be quite different. The agent does not act without keeping track of his actions; he cannot be intentionally self-controlling without being knowledgeably self-aware. In this way human agents are at once doers and perceivers, in a way that a divine agent is not.

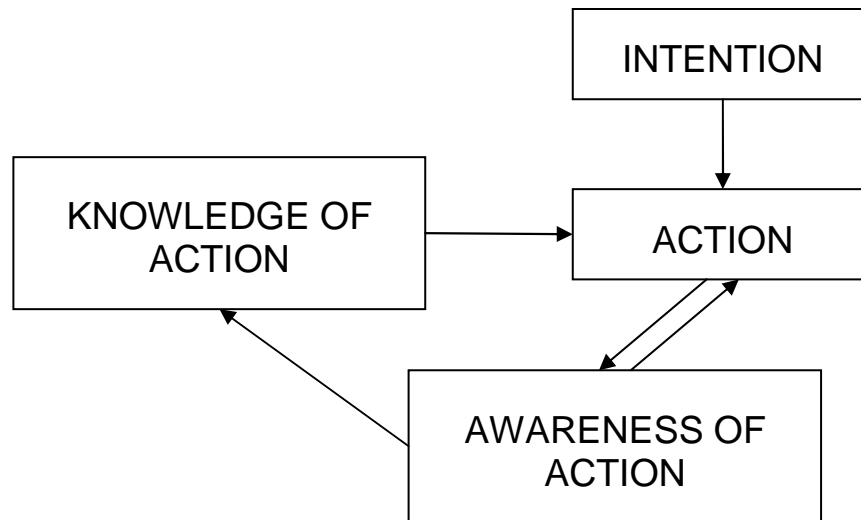


Figure 4: Non-observational, practical knowledge of one’s intentional actions.

This account of the non-observational, “practical” knowledge of one’s own intentional actions is illustrated very schematically in Figure 4. What it proposes is that the epistemic privilege characteristic of the knowledge of one’s own intentional actions is explained not by the fact that that knowledge has a special source, but instead by the special functional relationship that that knowledge bears to the action known: it is because one’s knowledge of what he is intentionally doing is essentially a *part* of the action itself that he knows about that action in a way that an outside observer cannot. And it is because of this integral relationship between action and knowledge of action that it does not make sense for an agent to have to “stop and look” to determine what he is intentionally doing: if this sort of thing were possible then the action in question would not be one of which the agent had a practical knowledge. This view preserves the first-personal/third-personal asymmetry that is fundamental to the knowledge of our intentional actions without denying that sense perception is implicated in the way we know what we are doing.

3.2 Intended and Unintended Actions

The difference between an agent’s knowledge of his intentional actions and the knowledge of those actions that can be had by an outside observer is not the only important asymmetry that Anscombe’s account of intention brings out; she also points us to an asymmetry within the first-

person perspective between the knowledge of intentional and *unintentional* actions, giving as an example a case in which “one noticed that one operated the traffic lights in crossing a road”, such that saying “I knew I was doing that, but only because I observed it” would be a way of marking the operation of the lights as unintentional (2000: 14). How does the notion of practical knowledge require us to think about this second asymmetry?

First, we should note that the accounts of non-observational knowledge that we criticized in Section 2 all have to analyze this example in terms of the idea that the knowledge of one’s intentional actions is essentially independent of sense perception; thus on these accounts it will be because the man has to *perceive* that he is operating the lights in order to know that he is that this counts as a case of knowledge by observation. But our discussion of the Self-Knowledge Problem showed us that such accounts get the metaphysics and epistemology of agency importantly wrong: for even the actions we engage in intentionally are not always known by us independent of our perceiving them to take place. If the division between intended and unintended actions parallels that between actions known by the agent in ways that depend and do not depend on observation, then this cannot be the right way to spell that latter distinction out.

Our alternative account of non-observational knowledge faces no such difficulty, however, as the claim that the man crossing the street has non-observational knowledge of his crossing but observational knowledge of his operating of the lights can be given quite a different sense. For in Anscombe’s example it is only the man’s knowledge *that he is crossing the street* that is “the cause of what it understands”, in the sense that that knowledge plays a role in keeping his behavior on course: if the man comes to know (or think) that he isn’t crossing the street after all but rather veering off into traffic, say, then he will try to adjust his behavior so as to do the thing he intends; while so long as he knows that he’s crossing the street, then all else being equal⁶ he will proceed contentedly along. By contrast the man’s knowledge that he is operating the traffic lights has no such status: it is knowledge that is entirely extrinsic to the behaviors in question, as the man has no stake in whether he operates the lights or not, and indeed if it seemed to us as if he *did* have such a stake in ensuring that he was operating the lights then we’d also find ourselves believing that that was not an unintended behavior after all. Thus it is, as Moran’s discussion of the example of the absent-minded walker suggests, possible to be aware of certain of one’s own actions from the perspective of observer rather than agent, and so to have a knowledge of them that is speculative rather than practical; it is just that when this is the case, the actions in question are ipso facto unintentional ones.

The crucial contrast, then, is not between empirical and non-empirical ways of knowing about an action but between knowledge that is causally implicated in shaping and sustaining the very action it is knowledge of, and knowledge that lacks this functional role. What the man crossing the street indicates when he says that he knows himself to be operating the lights “only because [he] observed it” is not that he relied on sense perception in coming to know this about himself, but rather that the knowledge in question is ultimately “accidental” to the action in the way that the knowledge of an intentional action never can be. From the perspective of the agent, that he has been operating the lights (or: casting sunlight on the wall, etc.) is something that he just “finds out” about himself in the same way that *we* might have discovered him to be doing that, while his knowledge that he was walking was quite different from this. With respect to his

⁶ “All else being equal” because it is possible to have multiple intentions, some of which can come into conflict with one another. For example, perhaps the man means to be crossing the street but also means *not* to be operating the lights, in which case if he knows that he’s succeeding in one of these intentions but not the other he will have to make a choice as to which is more important.

action under the former description but not the latter he simply “caught himself in the act”, we might say, and – unless perhaps he *then* decided that operating the traffic lights was something of a lark and so began to do it intentionally at that point – his knowledge of whether he was engaged in that action had nothing to do with whether the relevant events were brought about. This is just what it is to be an observer rather than an intentional agent, a passive taker-in of facts rather than a deliberate bringer-about.

3.3 Doing Without Knowing

I argued in Section 3.1 that in the ordinary case, what distinguishes an agent’s knowledge of his intentional actions as non-observational is that that knowledge is causally implicated in the very actions it is a knowledge of. Nevertheless, not *all* instances of human action proceed this way: we *ordinarily* act by keeping knowledgeable track of whether we are acting as we intend, but there are some cases where we are screened off from knowing about our actions in the ordinary way; and it is this kind of situation that can actualize a possibility that Anscombe explicitly rejects, namely that of an agent who acts intentionally in a given way but *lacks* the knowledge that he so acts. Anscombe makes several appeals to the supposed impossibility of such a situation at several different points in *Intention*, but for our purposes the most helpful one comes in one of her formulations of “practical knowledge”:

Practical knowledge is ‘the cause of what it understands’, unlike ‘speculative’ knowledge, which ‘is derived from the objects known’. This means more than that practical knowledge is observed to be a necessary condition of the production of various results; or that an idea of doing such-and-such in such-and-such ways is such a condition. It means that without it what happens does not come under the description—execution of intentions—whose characteristics we have been investigating. (2000: 87-88)

Here Anscombe is asserting that whenever an agent does not know that he is acting in a given way, any actions of his that are of that sort simply cannot be intentional ones; this is an idea she introduces much earlier on when she says that the reason-demanding “Why?” question appropriate only to intentional actions “is refused application by the answer: ‘I was not aware I was doing that’” (2000: 11). And unfortunately for Anscombe there are many possible cases of intentional action that have been adduced as counterexamples to this claim: a clerk may be making a stack of legible carbon copies just as he intends even while he is unsure whether he is actually pressing down hard enough (Davidson 1980: 50, 92); a man may think he is paralyzed and yet try and succeed to move his toes but without any awareness whether he does so; and so on.⁷ How should we understand agents’ epistemic relations to their intentional actions in (admittedly extraordinary) cases like these?

Given the view on offer here the most natural thing to say about such cases is that the actions in question are intentional insofar as the agents are *poised* to bring any awareness they may gain of what they are doing into the ordinary cycle of intentional self-control. Even in these conditions of relative ignorance, it remains that *if* the agents had known that they *hadn’t* been

⁷ I am grateful to Randall Amano for suggesting the second of these counterexamples, though a similar case has also been proposed by Kieran Setiya (2008: 290-91). Gibbons MS is a helpful discussion of what we should think about the significance of such cases for the possibility of privileged access to one’s own actions.

doing the things they intended to do, then they'd have been inclined to do something different to try to change that: thus Davidson's clerk is such that if he were made to *see* that he simply wasn't managing to produce all of the carbon copies he was after, then he'd begin pressing a bit harder or at least expressing some dissatisfaction with his failure to get the job done; and if we came to believe that although he was in fact producing the whole stack of carbon copies he didn't possess this sort of disposition toward succeeding then we'd no longer think of him as making all of the copies intentionally after all.

3.4 Two Knowledges?

Thus far in this section I have played up the ways in which the proposal now on offer is in keeping at least with the spirit of Anscombe's views on intentional action, but it is important to acknowledge the ways in which it diverges from the position staked out in *Intention*, too. For instance, here is how Anscombe resumes her analysis of the example we discussed at the start of Section 3.1:

Naturally my imaginary case, in which a man directs operations which he does not see and of which he gets no information, is a very improbable one. Normally someone doing or directing anything makes use of his senses, or of reports given him, the whole time: he will not go on to the next order, for example, until he knows that the preceding one has been executed, or, if he is the operator, his senses inform him of what is going on. This knowledge is of course always 'speculative' as opposed to 'practical'. Thus in any operation we really can speak of two knowledges—the account one could give of what one was doing, without averting to observation; and the account of exactly what is happening at a given moment (say) to the material one is working on. The one is practical, the other speculative. (2000: 88-89)

While these may be plausible things to say about the perceptual-cum-epistemic status of Anscombe's "director of operations", given the position developed here we can see that Anscombe is significantly misdescribing the role that sense perception plays in guiding and sustaining more common varieties of intentional action. For I have argued that, notwithstanding his usual reliance on what "his senses inform him of what is going on", the ordinary agent is *not* a purely "speculative" knower of what he intentionally brings about, is *not* experientially related "as observer" to those goings-on that constitute the material of his intentional doings, but rather implicates that perceptual knowledge in the distinctively "creative" manner that marks the activity of a self-knowing agent and makes that knowledge a cause of the very actions it understands. It is precisely for this reason that the sensory knowing of such an agent is not at all like, to borrow another of Anscombe's phrases, "a very queer and special sort of seeing eye in the middle of acting" (2001: 57); rather it is insofar as the eye of the agent is (as it were) in certain respects a "*doing* eye" that the agent can be a creator, and not just a passive observer of the events that take place around him.

Anscombe comes a bit closer to the view on offer here in another passage, where she discusses the role of sense perception in making successful action possible:

... I shut my eyes and write something. I can say what I am writing. And what I say I am writing will almost always in fact appear on the paper. Now here it is clear that my

capacity to say what is written is not derived from any observation. In practice of course what I write will very likely not go on being very legible if I don't use my eyes; but isn't the role of all our observation-knowledge in knowing what we are doing like the role of the eyes in producing successful writing? That is to say, once given that we have knowledge or opinion about the matter in which we perform intentional actions, our observation is merely an aid, as the eyes are an aid in writing. Someone without eyes may go on writing with a pen that has no more ink in it; or may not realise he is going over the edge of the paper on to the table or overwriting lines already written; here is where the eyes are useful; but the essential thing he does, namely to write such-and-such, is done without the eyes. So without the eyes he knows what he writes; but the eyes help to assure him that what he writes actually gets legibly written. In the face of this how can I say: *I do what happens*? If there are two ways of knowing there must be two different things known. (2001: 53)

As we have seen Anscombe does not respond to this last question by rejecting the idea that the "way of knowing" whose objects are the happenings in the world that the knower intentionally brings about is an observational way of knowing after all; instead, she seemingly accepts this equation, and works to preserve a sense in which our knowledge of our own actions is non-observational nevertheless. Yet in allowing that knowledge gained through ordinary sensory channels can be an "aid" in intentional action, she makes room to say something quite different, namely that it fills this role precisely by providing the agent with the sort of knowledge of what he does that is necessary for him to keep on doing it. For it is generally not enough, if one wishes to count as writing, just to stand over a paper and move a pen around: if certain sorts of things fail to be brought about then this may not end up being a case of writing at all, and so the empirical knowledge of his writing that an ordinary scribe will rely on to keep his letters legible will also be the source of the scribe's knowledge of "what he writes", i.e. his knowledge that he really is acting as he intends. On such a picture there are not two ways of knowing involved here after all, but rather only *one* way of knowing that has as its objects *both* the agent's intended actions *and* the happenings he intentionally brings about – those things being, of course, just two ways of describing what is really an indissoluble whole. Given the real possibility of failure, an agent often will not have non-empirical knowledge that he is doing what he intends; but the role of his sensory knowledge in bringing about the very action that it is knowledge of keeps the knowledge from becoming merely speculative.

Similar problems arise when we consider Anscombe's discussions of what Roger Teichmann (2009: 22) refers to as "Theophrastus' Principle", which marks the peculiar "direction of fit" had by agents' judgments of what they intend to do or be doing. Anscombe illustrates the idea with famous contrast between two sorts of lists:

Let us consider a man going round a town with a shopping list in his hand. Now it is clear that the relation of this list to the things he actually buys is one and the same whether his wife gave him the list or it is his own list; and that there is a different relation when a list is made by a detective following him about. If he made the list itself, it was an expression of intention; if his wife gave it to him, it has the role of an order. What then is the identical relation to what happens, in the order and the intention, which is not shared by the [detective's] record? It is precisely this: if the list and the things that the man buys do not agree, and if this and this alone constitutes a *mistake*, then the mistake is not in the list but in the man's performance (if his wife were to say: 'Look, it says butter and you have

bought margarine', he would hardly reply: 'What a mistake! we must put that right' and alter the word on the list to 'margarine'); whereas if the detective's record and what the man actually buys do not agree, then the mistake is in the record. (2001: 56)

The passage is complicated, but Anscombe's point should be clear. It would be foolish, for instance, to praise the list-maker's foresight if the shopper comes home with everything that was on his list; and by the same token it would make no sense to chide the man for simply buying everything the detective said he would. For the man was *supposed* to buy the things that were on the list, so if those things were available and he returns home without them then the fault is with him: the list is not a prediction of what the man will buy, but rather a guide for his purchasing. And the detective, for his part, writes down the things that the man buys *because the man buys them*: his list is not a prediction either, but nor does the man buy what he does because the detective has written it down. Thus if the detective turns out not to have written down all and only the things that the man bought, we say that his list is in error; whereas if there is a similar discrepancy between what the man buys and the list his wife sent him off with it is his *actions* that are in error, while the list is just fine as it is.

Yet Anscombe wants to extend the moral of this example much more broadly, and holds that *whenever* an agent makes a judgment as to what he is or intends to be doing, it can *only* be in his action that any mistake lies:

In some cases the facts are, so to speak, impugned for not being in accordance with the words, rather than *vice versa*. This is sometimes so when I change my mind; but another case of it occurs when e.g. I write something other than I think I am writing: as Theophrastus says (*Magna Moralia* 1189b 22), the mistake here is one of performance, not of judgment. There are other cases, too: for example St. Peter did not *change his mind* about denying Christ; and yet it would not be correct to say that he made a lying promise of faithfulness. (2001: 4-5)

And again:

... is there not another possible case in which a man is *simply* not doing what he says? As when I say to myself 'Now I press Button A'—pressing Button B—a thing which can certainly happen. This I will call the *direct* falsification of what I say. And here, to use Theophrastus' expression again, the mistake is not one of judgment but of performance. That is, we do *not* say: What you *said* was a mistake, because it was supposed to describe what you did and did not describe it, but: What you *did* was a mistake, because it was not in accordance with what you said. (2001: 57)

Anscombe is clearly right to insist that in cases like these, as Teichmann puts it, "if there is anything wrong with the [agent's] statement [of what he is intentionally doing], it is clearly not *simply* that the agent has misreported things" (2008: 24, emphasis added); but we should not follow her in holding that there is *no* sense in which these cases can involve mistaken judgments as well.⁸ For as we have seen, there are plenty of possible situations in which a person takes himself intentionally to be acting in a certain way and yet is failing to act as he thinks, and it is precisely for this reason that knowing whether we act as we intend will often require a perceptual

⁸ For a similar point see Moran 2000: 61.

attentiveness to those things we bring about. Thus what Anscombe and Teichmann ought to be saying is that in cases like these there can be faults that lie *both* in the judgment *and* in the performance; whereas when it comes to predictions, investigative reports, and other judgments that do not involve one's present intentional actions the same cannot be said. This fact is sufficient mark off a crucial aspect of the way in which an agent's self-knowledge is "practical", and we should not go so far as to say that its possessing this characteristic means that it cannot be in any way accountable to the facts.

4 Consciousness in Action

According to the view proposed in Section 3, our knowledge of our own intentional actions counts as non-observational not because it has a special source or lacks the dependence on sense perception characteristic of our knowledge of ordinary worldly happenings, but rather because of the distinctive functional role that an agent's perceptual awareness of his actions plays in his mental and behavioral economy: it is because an agent's self-awareness is caught up in a dynamic cycle of deliberate self-control that it cannot be understood as a form of mere self-observation. This view can seem to be threatened, though, by empirical research revealing certain sorts of dissociations between conscious visual experience and the visually-guided online control of action, research which has led some to hypothesize that these two capacities are subserved by distinct and largely independent visual subsystems, a "ventral" stream which supplies the stable representation of the world responsible for conscious vision and a more primitive "dorsal" stream comprised of egocentric representations dedicated to skillful visuomotor control. For present purposes, the crucial findings can be summarized as follows (see Goodale and Milner 2004: 17-30, 33-37, 82-89):

- Patients with various sorts of visual agnosias are unable to make accurate visually-based judgments but possess largely intact capacities for visually-guided behaviors such as reaching, grasping, walking, or posting a card through a slot whose orientation they are unable accurately to report.
- Patients with optic ataxia err frequently in reaching for or otherwise directing their behaviors toward physical objects in various spatial locations, yet can make accurate verbal descriptions concerning the very spatial properties with respect to which their motor behaviors are disrupted.
- Finally, even neurologically intact patients evince a distinction between conscious visual experience and "vision for action": for example, in the Ebbinghaus illusion (see Figure 5) the interior circles appear to have different sizes depending on the sizes of the circles that surround them, but when asked to reach and grasp the interior discs the scaling of your grasp will be unaffected.

Given a sufficiently stark articulation this "two visual systems" hypothesis makes untenable any view according to which conscious visual experience has a causal role to play in the online guidance of behavior, restricting it instead to supplying intentional directives and very coarse-grained behavioral specifications while leaving the dirty work of controlling behavior to the non-conscious midbrain stream; thus it can seem that the view of Section 3 is little more than a piece

of pre-scientific naïveté, an armchair philosophical myth that accords with many of our commonsense intuitions but stands in need of severe empirical correction.

Yet this stark articulation of the two visual systems hypothesis is by no means a perfect fit for the available empirical data, and a more nuanced understanding of their upshot is not nearly so problematic for the view on offer here.⁹ For one thing, subjects with severe dorsal impairments but intact conscious vision are not nearly as behaviorally inept as the stark articulation would predict: thus an optic ataxic will direct his actions toward an object more accurately, if not normally, when conscious visual feedback is allowed, revealing an ability to guide the structure of his movements by drawing on what is consciously seen (Jeannerod 1997: 57); and more generally the visuomotor deficits associated with optic ataxia do not show up when actions are targeted at stationary objects in the center of the field of vision (Rossetti et al 2003). This suggests that even in the case of relatively simple intentional behaviors, the agent's conscious visual awareness of how he is moving his body stands poised to play a role in the control of action, even if normally this is not required.




Figure 5: The Ebbinghaus illusion. The interior circles are the same size in each array, yet the one on the right-hand side appears larger because of the respective sizes of the circles that surround each.

Even more significantly for our purposes, the stark articulation contradicts the finding that subjects with intact dorsal systems and impaired conscious vision are far from behaviorally normal: for example, Milner and Goodale's visual agnostic patient "DF" can engage successfully in certain sorts of visuomotor tasks but is unable to structure her reaching and grasping in ways appropriate to the use or function of the objects she targets. As Jeannerod puts it, the visuomotor function remaining in subjects with severe ventral system deficits "is a very impoverished one":

It corresponds to the activity of subcortical projections to parietal areas which bypass V1, as part of a 'primitive' system for fast and crude reactions to visual stimuli ... It might be that this system operates only in situations of emergency, automatic movements or decreased awareness. ... In fact, the type of function which is normally carried out by parietal structures is more than a simple, more or less direct, visuomotor transformation

⁹ The argument of the next two paragraphs is especially indebted to a number of conversations with John Campbell, as well as the discussion of the two visual systems hypothesis in Briscoe 2009 and an unpublished manuscript of Alva Noë's.

... Parietal function does not operate in isolation, [but] is embedded in a broader system for producing action, which involves other areas, including those from the ventral system. (1997: 71-72)

The stark articulation of the two visual systems hypothesis rejects this sort of nuance, instead likening human action to the behavior of a teleassistant robot (Goodale and Humphrey 1998: 202) or a heat-seeking missile (Campbell 2002: 56), with conscious experience figuring only when it comes time to establish the parameters and non-conscious, automatic processing then stepping in to complete the task that has been set for it. But the more nuanced view, which appears to be *at least* as compatible with the available experimental data as the starker one is, holds by contrast that “this dichotomous thinking was an oversimplification”, and that “there are several routes for action, rather than separate routes for action and perception” (Jeannerod 1997: 81-82). And it is only the stark articulation, not the nuanced one, that is incompatible with the idea that perceptual experience provides an agent with a knowledge of his actions that functions in turn as *causa rerum intellectarum*, the cause of the very things it understands.

Bibliography

- Acredolo, L. 1990. Behavioral approaches to spatial orientation in infancy. *Annals of the New York Academy of Science* 608: 596-612.
- Alexander, P. 1985. *Ideas, Qualities, and Corpuscles: Locke and Boyle on the External World*. New York: Cambridge University Press.
- Anscombe, G. E. M. 2000. *Intention*, 2nd ed. Cambridge: Harvard University Press.
- Aquinas, T. 1983. *Treatise on Happiness*. Trans. J. A. Oesterle. Notre Dame, IN: University of Notre Dame Press.
- Avery, T. 2009. Vision and the first person. Presented at the Berkeley-London Conference, 5/8/2009.
- Bermúdez, J. L. 1995a. Transcendental arguments and psychology: the example of O'Shaughnessy on intentional action. *Metaphilosophy* 26: 379-91.
- Bermúdez, J. L. 1995b. Ecological perception and the notion of a nonconceptual point of view. In *The Body and the Self*, ed. J. L. Bermúdez, A. Marcel, and N. Eilan, 153-173. Cambridge: The MIT Press.
- Bermúdez, J. L. 1998. *The Paradox of Self-Consciousness*. Cambridge: The MIT Press.
- Briscoe, R. 2009. Egocentric spatial representation in action and perception. *Philosophy and Phenomenological Research* 79: 423-60.
- Campbell, J. J. 1993. *Past, Space, and Self*. Cambridge: The MIT Press.
- Campbell, J. J. 2002. *Reference and Consciousness*. New York: Oxford University Press.
- Campbell, J. 2007. What's the role of spatial awareness in visual perception of objects? *Mind and Language* 22: 548-62.
- Campbell, J. 2008. Consciousness and reference. In *The Oxford Handbook of Philosophy of Mind*, ed. B. McLaughlin and A. Beckermann, 648-62. New York: Oxford University Press.
- Cassam, Q. 1994. *Self and World*. New York: Oxford University Press.
- Cassam, Q. 2007. *The Possibility of Knowledge*. New York: Oxford University Press.
- Davidson, D. 1980. *Essays on Actions and Events*. New York: Oxford University Press.
- Donnellan, K. 1963. Knowing what I am doing. *Journal of Philosophy* 60: 401-09.
- Evans, G. 1982. *The Varieties of Reference*. Ed. J. McDowell. New York: Oxford University Press.
- Falvey, K. 2000. Knowledge in intention. *Philosophical Studies* 99: 21-44.
- Freidman-Hill, S., L. C. Robertson, R. Desimone, and L. G. Ungerleider. 2003. Posterior parietal cortex and the filtering of distractors. *Proceedings of the National Academy of Sciences* 100: 4263-68.
- Friedman-Hill, S., L. Robertson, and A. Treisman. 1995. Parietal contributions to visual feature-binding: evidence from a patient with bilateral lesions. *Science* 269: 853-55.

- Gallagher, S. 2005. *How the Body Shapes the Mind*. New York: Oxford University Press.
- Gallagher, S., and D. Zahavi. 2008. *The Phenomenological Mind: An Introduction to Philosophy of Mind and Cognitive Science*. New York: Routledge.
- Gibbons, J. MS. Seeing what you're doing. Forthcoming in *Oxford Studies in Epistemology*.
- Gibson, J. J. 1986. *The Ecological Approach to Visual Perception*. Hillsdale: Lawrence Erlbaum & Associates.
- Goodale, M. A. and G. K. Humphrey. 1998. The objects of action and perception. *Cognition* 67: 181-207.
- Goodale, M. A. and A. D. Milner. 2004. *Sight Unseen: An Exploration of Conscious and Unconscious Vision*. New York: Oxford University Press.
- Harman, G. 1990. The intrinsic quality of experience. *Philosophical Perspectives* 4: 31-52.
- Heidegger, M. 1964. *Being and Time*. Trans. J. Macquarrie and E. Robinson. New York: Harper & Row.
- Husserl, E. 1970. *The Crisis of the European Sciences and Transcendental Phenomenology*. Trans. D. Carr. Evanston: Northwestern University Press.
- Husserl, E. 1989. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, Second Book: Studies in the Phenomenology of Constitution*. Trans. R. Rojcewicz and A. Schuwer. Dordrecht: Kluwer.
- Husserl, E. 1997. *Thing and Space: Lectures of 1907*. Trans. R. Rojcewicz. Dordrecht: Kluwer.
- Jeannerod, M. 1997. *The Cognitive Neuroscience of Action*. Blackwood, N.J.: Blackwell.
- Kant, I. 1933. *Critique of Pure Reason*. Trans. N. Kemp Smith. New York: Palgrave Macmillan.
- Kant, I. 1997. *Prolegomena to Any Future Metaphysics*. Trans. G. Hatfield. New York: Cambridge University Press.
- Lishman, J. R., and D. N. Lee. 1973. The autonomy of visual kinaesthesia. *Perception* 2: 287-94.
- Millikan, R.G. 2002. *The Varieties of Meaning*. Cambridge: The MIT Press.
- Moran, R. 2000. Anscombe on 'practical knowledge'. In *Agency and Action*, ed. J. Hyman and H. Steward, 43-68. New York: Cambridge University Press.
- Noë, A. 2004. *Action in Perception*. Cambridge: The MIT Press.
- O'Shaughnessy, B. 1980. *The Will: A Dual Aspect Theory*, vol. 2. New York: Cambridge University Press.
- Palmer, S. E. 1999. *Vision Science: From Photons to Phenomenology*. Cambridge: The MIT Press.
- Paul, S. K. MS. How we know what we're doing. Forthcoming in *Philosophers' Imprint*.
- Peacocke, C. 1983. *Sense and Content: Experience, Thought, and Their Relations*. New York: Oxford University Press.
- Perry, J. 1993. Thought without representation. In his *The Problem of the Essential Indexical and Other Essays*, 205-26. New York: Oxford University Press.

- Robertson, L. 2004. *Space, Objects, Minds, and Brains*. New York: Psychology Press.
- Robertson, L. and A. Treisman. 2006. Attending to space within and between objects: implications from a patient with Balint's syndrome. *Cognitive Neuropsychology* 23: 448-62.
- Robertson, L., A. Treisman, S. Friedman-Hill, and M. Grabowecky. 1997. The interaction of spatial and object pathways: evidence from Balint's syndrome. *Journal of Cognitive Neuroscience* 9: 295-317.
- Rock, I. 1983. *The Logic of Perception*. Cambridge: The MIT Press.
- Roessler, J. 2003. Intentional action and self-awareness. In *Agency and Self-Awareness*, ed. J. Roessler and N. Eilan, 382-405. New York: Oxford University Press.
- Rossetti, Y., L. Pisella, and A. Vighetto. 2003. Optic ataxia revisited: visually guided action versus immediate visuomotor control. *Experimental Brain Research* 153: 171-79.
- Sartre, J. P. 1992. *Being and Nothingness*. Trans H. Barnes. New York: Washington Square Press.
- Scheer, J. 2009. Experience and self-consciousness. *Philosophical Studies* 144: 95-105.
- Setiya, K. 2008. Practical knowledge. *Ethics* 118: 388-409.
- Shams, L., Y. Kamitani, and S. Shimojo. 2002. Visual illusion induced by sound. *Cognitive Brain Research* 14: 147-52.
- Strawson, P. 1966. *The Bounds of Sense*. London: Methuen.
- Stroop, J. R. 1935. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* 18: 643-62.
- Treisman, A. 1988. Features and objects: the fourteenth Bartlett Memorial Lecture. *The Quarterly Journal of Experimental Psychology Section A* 40: 201-37.
- Treisman, A. 1996. The binding problem. *Current Opinion in Neurobiology* 6: 171-78.
- Treisman, A. and G. Gelade. 1980. A feature-integration theory of attention. *Cognitive Psychology* 12: 97-136.
- Velleman, D. 1989. *Practical Reflection*. Princeton: Princeton University Press.
- Warren, D. 1998. Kant and the apriority of space. *Philosophical Review* 107: 179-224.
- Wittgenstein, L. 1975. *Philosophical Remarks*. Trans. R. Hargreaves and R. White. Chicago: University of Chicago Press.