WASHINGTON UNIVERSITY

Philosophy-Neuroscience-Psychology Program

Department of Philosophy

Dissertation Committee:

William Bechtel, Chairperson
Andy Clark
Jesse Prinz
Rick Grush

OBJECTIVE SUBJECTIVITY: ALLOCENTRIC AND EGOCENTRIC REPRESENTATIONS IN THOUGHT AND EXPERIENCE

by

Peter John Mandik

A dissertation presented to the
Graduate School of Arts and Sciences
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

May 2000

Saint Louis, Missouri

Contents

List of Figures

Acknowledgments

Portions of chapters 1 and 4 appeared in "Objectivity without space" in *The Electronic Journal of Analytic Philosophy, Special Issue on the Philosophy of Gareth Evans*. Ed. Rick Grush (1998) (http://www.phil.indiana.edu/ejap/1998/mandik98.html). The second section of chapter 2 is from "Qualia, space, and control" in *Philosophical Psychology* 12 (1): 47-60 (1999). The remaining sections of chapter 2 and a version of chapter 3 will appear as "Points of view from the brain's eye view: Subjectivity and neural representation" in W. Bechtel, P. Mandik, J. Mundale, and R. Stufflebeam (eds.), *Philosophy and the Neurosciences: A Reader* (Blackwell).

In developing this work I have incurred numerous debts, intellectual and other. I am deeply grateful for the unparalleled intellectual stimulation I received in the midst of my fellow graduate students at Washington University both in the Philosophy-Neuroscience-Psychology Program and in the so-called "regular" Philosophy Program. In particular I thank Tad Zawidzki, Chase Wrenn, Chris Eliasmith, Whit Schonbein, and Jon Waskan for being there. Nowhere else have I encountered such an excellent collection of minds with which to think and drink. You guys rock.

For igniting much of my interest in the core topics explored in this dissertation, I thank Rick Grush. Rick is not only a collaborator and friend, but a seemingly unending

source of favors and opportunities that I hope to some day be able to repay. Spending time with him in Denmark at the University of Aarhus in 1998 and at the University of Pittsburgh in 1999 were wonderful times in which I basked in Rick's unapologetic quest for truth, excellence, and kicking butt.

To Bill Bechtel—friend, collaborator, mentor, and adviser—I am eternally grateful. Bill is yet another source of inspiration and opportunity I hope to someday be able to repay. Among other things, Bill is my paradigm for how to be a productive participant in the philosophical community.

For discussion of the material in this dissertation, I am grateful to the aforementioned individuals as well as Jesse Prinz, Brian Keeley, Andy Clark, Pepa Toribio, Roger Gibson, Eric Brown, Peter Machamer, Joe Ullian, George Graham, William Wimsatt, and Dan Weiskopf.

Thanks to my parents Peter and Sandra Mandik and my brother Michael Mandik for being the best family I could ever hope for.

For love, support, friendship, and everything else, I will never stop thanking my wife, Ann Mandik.

Overview of the Chapters

One common use of the notions of objectivity and subjectivity is to demarcate kinds of judgment (or thought or belief). On such a usage, prototypically objective judgments concern matters of empirical and mathematical fact such *as the moon has no atmosphere* and *two and two are four*. In contrast, prototypically subjective judgments concern matters of value and preference such as *Mozart is better than Bach* and *anchovy ice cream is disgusting*. In chapter 1, I examine various theoretical proposals for how to account for the objectivity and subjectivity of judgments. The theory I argue for—a correspondence theory of objectivity—explicates the objectivity/subjectivity of a judgment in terms of the mind-independence/mind-dependence of the properties picked out by the predicates in the judgment. Among the proposals that I argue against are the suggestions that the objectivity of a judgment requires that the judgment is agreed upon (the consensus theory) or that the judgment not have indexical constituents (the indexical theory).

Central to my arguments in this chapter is the question of what the bearers of the properties of subjectivity and objectivity are. I consider three candidates. The first holds that entire judgments are the bearers of objectivity or subjectivity. The second and third candidates hold that the bearers are the conceptual constituents of judgments, that is, conceptual equivalents of general and singular terms respectively. I argue for the view

that general (predicative) concepts (construed as mental representations) are the ultimate bearers of objectivity and subjectivity. I then argue that this result is best explained by a correspondence theory, for indexical theories locate objectivity and subjectivity at singular term constituents of judgment and consensus theories treat entire judgments as the bearers of objectivity and subjectivity.

In chapter 2 I turn to further develop the correspondence theory of objectivity against a backdrop of naturalized accounts of mental representation. Using examples drawn from cognitive scientific research on visual object recognition, navigation, and mental imagery, I develop a notion of subjectivity modeled on the point-of-view or perspective embodied in imagistic representations. To see how pictorial representations embody a point of view consider two photographs of the same object taken from two different angles. Compare, for example, two photographs taken of a person's face: The first may be head-on, the other may show the head in profile. The camera that produced the photos occupied two different points of view with respect to the person's head. The different positions and orientations of the camera constitute the points of view of the camera with respect to the subject. The representational contents of the photographs produced include content about these points of view. We can tell by looking at the photographs whether the camera was in front of or to the side of the person's face. I offer that the distinction between subjective and objective (egocentric and allocentric) representations is a distinction between representations that have and representations that lack this kind of point of view. The essence of this kind of point of view is a property of

mental representations that include in their content the representing subject. Spelling this out further yields the following definitions. A subject *S* has a subjective representation *R* of *X* if and only if the representational content of *R* includes relations *S* bears to *X*. A subject *S* has an objective representation *R* of *X* if and only if the representational content of *R* does not include relations that *S* bears to *X*.

In this chapter I also sketch a way of understanding the objective/subjective distinction within naturalized accounts of representation that construe representational content as being determined by causal relations between representations and what they represent. I show how the above definitions may be spelled out within a framework that explains *R*'s representing *X* in terms of *R*'s having the function of carrying information about *X*. I show how these kinds of approaches fit with the correspondence theory by discussing cases in which what is represented exists independently (or dependently) of the representing subject.

In chapter 3 I take up the topic of the subjectivity of conscious experience. Central to my discussion is the infamous Nagel-Jackson 'knowledge argument' and the indexical theory of subjectivity that has grown largely as a response to it. I show how the correspondence account is able to deal with the considerations raised by the knowledge argument without falling prey to problems that befall indexical theories.

The gist of the knowledge argument is as follows. A person that has never had any experiences as of seeing a red thing may nonetheless have exhaustive knowledge of the physical goings on in the nervous system of an individual seeing red. Such a knowledgeable person may know all the physical facts about seeing red without having a red experience. But suppose this knowledgeable individual was to finally have a red experience. Many find it intuitive to suppose that such an individual would learn something new, namely, they would learn what it is like to see red. Anti-physicalistic conclusions are supposed to follow on the supposition that learning what it is like to see red involves learning some new fact. Prior to having the experience, the subject knew all the physical facts; thus in learning a new fact after having a red experience, the subject learns a non-physical fact. Thus, allegedly, physical facts do not exhaust all the facts, since they do not include certain facts about experience. Knowledge of physical facts leaves out knowledge of subjective facts.

One kind of physicalist response to the knowledge argument has been to grant that something new (and subjective) is learned, while giving a physicalistic account of what subjectivity amounts to. Foremost among such responses is the indexical response due to philosophers such as Lycan (1987, 1996) Tye (1995) and Rey (1997). To take one instance as representative: Lycan's account of subjectivity is as follows. Experiences are representations. My visual experience of my blue coffee mug is a mental representation of the mug as being blue. When I introspect my experience, I form a second-order representation of the first-order representation of the coffee mug. Other people may form

syntactically similar second-order representations, but those representations will be about their first-order states, not mine. The crucial analogy here is to the use of indexicals in speech. When I say "my leg hurts" I am referring to my leg, and only I can refer to my leg by using that utterance. You may use a syntactically, morphologically, and phonologically similar construction: you may utter the words "my leg hurts", but in doing so, you would be representing your leg, not mine. Analogously, only I can represent my first-order states by the introspective application of self-referential indexical concepts. And this, according to Lycan, is the ultimate explication of subjectivity.

One pressing problem for the indexical account is whether the indexical/nonindexical distinction can be accounted for in naturalistic, yet alone neuroscientific, terms. A prevalent characterization of indexicality is that it arises when the content of a representation depends on context. Many naturalistic accounts of representation are externalist, in that they favor a view of content whereby a representation has its content in virtue of certain relations it bears to environmental states. If the representation bore different relations to different environmental states, then it would have different content. The problem is that such an externalist account of representation makes all representation look indexical, thus there is no indexical/nonindexical distinction with which to reconstruct a subjective/objective distinction. I propose to circumvent this and other problems by showing how the correspondence theory is able to handle the considerations raised by the knowledge argument at least as well as indexical theories of subjectivity.

In chapter 4 I examine several ways in which the topic of spatial representation bears on the offered account of objectivity/subjectivity. There is a long tradition of discerning deep connections between the notions of space, objectivity and subjectivity. Much of the importance of the this chapter is to reevaluate the alleged connections. I address arguments due to Peter Strawson and Gareth Evans that space is a necessary requirement for objectivity. I find their arguments wanting and argue instead that on several interpretations of their thesis, it is false.

Though my project is largely concerned with metaphysical as opposed to epistemological questions ("Are all *X*s also *Y*s?" as opposed to "How do you know all *X*s are also *Y*s?") I do not wish to remain silent on epistemological questions. I turn to such questions in chapter 5. I address the question of how one comes to know whether something is objective or subjective. The answer to this question is given by way of showing how the correspondence theory accounts for the traditional association of the objective with the public and the subjective with the private. This account grows directly out of my explication of metaphysical objectivity and subjectivity.

Chapter 1: Subjective and Objective Judgments

1.0 Introduction

Many philosophical issues concern questions of objectivity and subjectivity. Of these questions, there are two kinds. The first considers whether something is objective or subjective; the second what it *means* for something to be objective or subjective—questions that inquire as to the very essence of objectivity and subjectivity. I call questions of the first kind "questions of application" and questions of the second kind "questions of constitution".

Examples of questions of application include (but are not limited to) the following. Is it possible for science to be objective, or are the claims of science inextricably bound up with subjective points of view? If science is indeed objective, is there anything that escapes its grasp in virtue of being essentially subjective (as some have claimed for consciousness)? Are the things that comprise the subject matter of scientific and folk theories things that exist objectively or are they instead the subjective results of the way our ways of thinking and talking carve up the world? Must all physical things exist objectively and vice versa?

Examples of questions of constitution include (but are not limited to) the following. Is objectivity an unachievable ideal that can only be approximated by degrees of intersubjectivity? How deep is the common analogy between objectivity and subjectivity on the one hand and the literary conventions of third person and first person points of view on the other? Are subjectivity and objectivity ways that things exist (mind-dependently vs. mind-independently) or are they ways of representing things (from a particular perspective or point of view vs. from no particular perspective or point of view)? If subjectivity and objectivity are ways of representing, then what are the proper roles of notions of truth and indexicality in a theory of objectivity and subjectivity?

In this dissertation I address both questions of application and questions of constitution. The primary aim of this chapter is to begin by addressing the question of what it means for judgments to be objective or subjective. One common use of the notions of objectivity and subjectivity is to demarcate kinds of judgment (or thought or belief). On such a usage, prototypically objective judgments concern matters of empirical and mathematical fact such as *the moon has no atmosphere* and *two and two are four*. In contrast, prototypically subjective judgments concern matters of value and preference such as *Mozart is better than Bach* and *vanilla ice cream with ketchup is disgusting*. I offer these examples not to take sides on whether such judgments actually are objective or subjective, but only to call attention to a typical way of using "objective" and "subjective". The question arises as to what it means in this context to call these respective judgments "objective" and "subjective". Some have proposed that the

difference hinges on truth. Objective judgments are absolutely true, whereas the truth of subjective judgments is relative to the person making the judgment: my judgments are true for me, your judgments are true for you. You and I can each utter "vanilla tastes great" but in your mouth this may constitute a truth and in my mouth it may constitute a falsehood. Subjective judgments are subject relative. Some philosophers have noted an analogy between this kind of subject relativity and a kind that obtains for indexical expressions. You and I can both utter "I am here" and thereby express different propositions. Some philosophers have construed indexicality as an instance of subjectivity and some others have even gone so far as to argue that subjectivity just is indexicality.

I will postpone taking sides on these issues, but let me spell out further what I take the importance of the above remarks to be. I call attention to the practice of labeling judgments (and beliefs etc.) objective and subjective. In this discussion, it is representations that have propositional or sentential structure that are the first and foremost instances of objective (and subjective) things. The question arises, then, of what it is about these representations that makes them subjective. One suggestion is that the subjective/objective distinction marks a distinction in ways of assigning truth values to these representations, ways that are relativist and absolutist, respectively. Another suggestion is that the subjective/objective distinction marks a distinction in ways of assigning representational content to these representations, ways that are indexical and non-indexical, respectively. Yet another approach seeks to classify representational

schemes in terms of the degree to which they reflect a particular perspective or point of view in the literal sense that pictorial representations represent the visual appearance of objects from a point of view. On this suggestion, pictures are the prototypically subjective representations and objective representations are to be defined in contrast. Among the issues to be sorted out in considering the "truth", "indexical", and "picture" suggestions are those concerning whether they constitute distinct viable alternatives, and if so, whether they are compatible. Such sorting will have to wait. Ultimately, however, I advocate a combination of the "truth" and "picture" approaches.

1.1 Remarks on Methodology

"Objectivity" and "subjectivity," it goes without saying, are words. It perhaps does not go without saying that these words are semi-technical. "Objectivity" and "subjectivity" have multiple overlapping meanings in the mouths of professional philosophers as well as in the mouths of both non-professionals and non-philosophers. The question arises as to what these words mean. Had I been writing a few decades ago, I would proclaim without apology or qualification that the concepts of objectivity and subjectivity were in need of analyses. In the present, however, qualification and apology are in high demand. Conceptual analysis has fallen on hard times. Quinean worries about the analytic/synthetic distinction loom large. Also, successful analyses are just too few and far between to maintain enthusiasm for such projects. If an analysis is (i) a

specification of the necessary and sufficient conditions for the correct application of a term and (ii) captures all of the pretheoretic intuitions regarding the meaning of the term, then there are embarrassingly few analyses. (Even the old standby that *bachelor* analyses into *unmarried male* has its counter-examples: the pope is an unmarried male, but it seems unintuitive to proclaim him a bachelor.) However, successor subjects to conceptual analysis have arisen. I have mind two: What Quine (1960, pp. 258-260) calls "explication" and what Rawls (1993, p. 8) calls "reflective equilibrium".

Quinean explication is at heart definition by stipulation. The stipulation earns its keep by being both theoretically useful and having some (but not necessarily total) semantic overlap with its pre-theoretic homophonic forbear. Quine writes:

> We do not claim synonymy. We do not claim to make clear and explicit what the users of the unclear expression had unconsciously in mind all along. We do not expose hidden meanings, as the words 'analysis' and 'explication' would suggest; we supply lacks. We fix on the particular functions of the unclear expression that make it worth troubling about, and then devise a substitute, clear and couched in terms to our liking, that fills those functions. Beyond those conditions of partial agreement, dictated by our interests and purposes, any traits of the explicans come under the head of "don't-cares". . . .Under this head we are free to allow the explicans all manner of novel connotations never associated with the explicandum. (1960, pp. 258-259)

According to Rawls' reflective equilibrium method, a moral principle that cannot be deduced from any other principle can be justified by seeing if observing the principle would yield the considered judgments of competent judges. If it does not, we have two choices: we can adopt a different principle or change the considered judgments. The

process of adjusting the principles and the judgments so that they are brought into accord is the method of reflective equilibrium. Any first principles arrived upon by this method count as justified, even though they were not deducible from any other principles. Rawls' notion of reflective equilibrium need not apply only to moral matters. Viewing concepts as principles regulating the applications of terms allows viewing reflective equilibrium as a successor to conceptual analysis. Reflective equilibrium as a successor to conceptual analysis involves comparing a proposed definition of a term to intuitions regarding proper use and evidence regarding theoretical utility. If a definition is theoretically useful and comports with many pretheoretic intuitions, then residual pretheoretic intuitions that cannot be accounted for may simply be disregarded. Viewed this way, there is little if any difference between the methods of reflective equilibrium and Quinean explication.

Returning to the question of the meaning of "objectivity" and "subjectivity", what I am after, then, is not so much analyses of those terms as explications. My method will involve reflective equilibrium. My explications will earn their keep by doing theoretical work in the naturalized philosophy of mind. Professional philosophers of mind have used "objectivity" and "subjectivity" both technically and non-technically. Some have offered theories and definitions. What follows in this chapter and others are definitions and applications of the terms that do more work than their theoretical and pre-theoretical precursors.

1.2 Epistemic Objectivity and Metaphysical Objectivity

In this section I discuss two senses of "objectivity" (and "subjectivity"): what I shall call the "epistemic" sense and the "metaphysical" sense, respectively. Various philosophers have distinguished these two senses (See, for example, Audi 1992; Bell 1992; Foss 1993; Newell 1986; Searle 1992: 93-100; Rescher 1997: 3-5). I begin by discussing the epistemic sense of the objective/subjective distinction.

1.2.1 Introducing Epistemic Objectivity and Subjectivity

Many people, philosophers and non-philosophers alike, hold the following view of matters of taste. Judgments of taste are subjective. If you judge the Mona Lisa to be beautiful, and I judge the Mona Lisa to be ugly, we can both be right. Such a view of matters of taste is oft expressed by the slogan "beauty is in the eye of the beholder". I mention this view of matters of taste not to assert its truth or falsity, but because its familiarity helps to focus attention on a common use of the terms "objective" and "subjective". On this use, objectivity and subjectivity are properties of judgments. Another familiar view is the view that judgments in mathematics and the natural sciences are objective. If you judge a sample of gold to weigh five grams and I judge the same sample to weigh six grams, then at least one of us is wrong. Unlike our differing judgments of the beauty of the Mona Lisa which can both be right, our differing

judgments of the mass of a sample of gold cannot both be right. The same sort of thing is held to apply to mathematical judgments. If you judge the sum of two numbers to be 1,082 and I judge the sum to be 1, 083, then at least one of us has to be wrong. If we disagree on the result of adding two numbers, we cannot both be right. I mention the view that judgments of mathematics and natural sciences are objective not to take a stand on the view, but instead to direct the reader's attention to a familiar use of the terms "objective" and "subjective".

Calling judgments objective or subjective as in the above examples is to use what I shall call the "epistemic" senses of the terms "objective" and "subjective"—it is to speak of epistemic objectivity and epistemic subjectivity. This is the kind of objectivity and subjectivity had by judgments and things that are similar to judgments in the following respects. Judgments are the sorts of things that are either true or false, accurate or inaccurate. Like beliefs, percepts, news reports, books, articles, sentences, and pictures, they represent or portray things. They are about things—they have aboutness or intentionality. Only things that are representational can be epistemically objective or subjective. The sorts of things that can be true or false can be epistemically objective or subjective. There is thus a close connection between truth and falsity on the one hand and epistemic objectivity and subjectivity on the other. But I must stress that this close connection is not one of identity. In the above examples concerning the Mona Lisa and the sample of gold the distinction between objectivity and subjectivity is not the distinction between truth and falsity. To call the judgment that the Mona Lisa is beautiful

an epistemically subjective judgment is not to consign it to falsity, but instead to allow that it and its negation can both be true. In this case, the truth of the judgments depends on who made them. To call the judgment that the sample of gold weighs six grams objective is not to assert the truth of the judgment, for both the judgment that the gold weighs six grams and its negation can be objective even though they cannot both be true.

I turn now to discuss briefly three kinds of theories of epistemic objectivity: consensus theories, indexical theories, and correspondence theories. I eventually defend a correspondence theory of objectivity.

Consensus theories define the epistemic objectivity of a judgment in terms of agreement. Christopher Gauker (1995) defends a consensus theory whereby the objectivity of a judgment has to do with the likelihood that it would be agreed upon if rational people were to entertain the judgment. More specifically, Gauker construes the degree to which a judgment is objective as a function of three dimensions: the probability agents would come to agree, the degree of rationality of the agents, and the level of agreement (1995, p. 173).

Indexical theories capitalize on the naturalness with which we describe the subjective/objective distinction in terms of the first- and third- person points of view. Indexical theorists include William Lycan (1996), Georges Rey (1997), Michael Tye (1995), and A.W. Moore (1997). The gist of the theory is as follows. William Lycan's

biography, as told in the first-person point of view, would be chock-full of statements employing self-referential indexicals. Examples might include “I am a philosopher” and “I have a mustache”. If told instead from the third-person point of view Lycan’s biography would instead contain statements devoid of indexicals. Examples might include “William Lycan is a philosopher” and “William Lycan has a mustache”. In indexical theories of subjectivity, the subjectivity of a representation is explicated by the representation’s having indexical components. A natural suggestion, then, is that representations that lack indexical components are objective. An indexical account of epistemic subjectivity and objectivity, then, would define as subjective representational states that had indexical components and define as subjective representational states that lacked indexical components.

Its worth noting, however, that while Lycan explicates subjectivity in terms of representations that have indexical components, it is not clear that he intends to explicate objectivity in terms of representations lacking indexical components. Indeed, sometimes when he uses the word “objective”, he is referring to things that are not representations, indicating that while he is using the word “subjective” in its epistemic sense he is using the word “objective” in its metaphysical sense. It may not be accurate then, to describe Lycan as an indexical theorist of epistemic subjectivity and objectivity, since it is not clear what his views on epistemic objectivity are. I will continue, however, to classify his view on subjectivity as epistemic.

The core notion of correspondence theories is that an epistemically objective belief or sentence must be in some sense about something metaphysically objective. Rorty (1997) describes this notion of objectivity as a "mirroring" notion of objectivity for it involves the notion of representing things as they really are, that is, the way they are independent of the way they are represented. A description of something as being a hunk of titanium would be epistemically objective, according to the correspondence theory, because something can be a hunk of titanium independently of any one's representing it as such. In contrast, my belief that Brussels sprouts are disgusting is epistemically subjective because being disgusting requires being mentally represented as disgusting. To contort a cliché: Being disgusting is in the mouth of the taster.

Thus ends my initial discussion of epistemic objectivity. Of course, more can be said, and since my aim in this dissertation is to articulate a theory of epistemic objectivity, more will be said. But I satisfy myself here with conveying only a flavor of epistemic objectivity. The key points in the above discussion are (i) that only things with representational content can be either epistemically objective or epistemically subjective and (ii) the distinction between epistemic objectivity and subjectivity is orthogonal (though not entirely unrelated) to the distinction between truth and falsity. I turn now to discuss the metaphysical sense of the objective/subjective distinction.

1.2.2 Introducing Metaphysical Objectivity and Subjectivity

At the heart of the notion of metaphysical objectivity is the notion mind-independent existence. Metaphysically objective things are those that do not depend on minds for their existence. In contrast, metaphysically subjective things do depend on minds for their existence.

The metaphysical distinction between objectivity and subjectivity may be conveyed in several philosophical examples. Locke's distinction between primary and secondary properties is a distinction between metaphysically objective and subjective properties, respectively. Primary properties, for instance, properties like shape and mass were said to be instantiated by objects whether or not anyone had experienced or would experience those objects as having shape or mass. In contrast, secondary properties like color and odor were held to be properties of objects specifiable only by reference by the causal powers they exerted over human perceivers. An apple instantiates the property of being red only insofar as it has a disposition to cause a certain class of perceivers to see it as red. Redness, so the story goes, does not genuinely inhere in objects the way that roundness does. Red things depend on the minds of human perceivers for their redness, but round things would be round even though there were no minds to perceive them as such.

Roundness and redness, on the Lockean story, are metaphysically objective and subjective, respectively. Berkeley famously rejected this part of Locke's story by arguing that no distinction between primary and secondary properties can be drawn. Berkeley argued that so-called primary properties were just as perceiver dependent as so-called secondary properties. Thus Berkeley defended the metaphysical position of Idealism, the view that everything that exists either is a mind or depends for its existence on being perceived (or at least thought about) by minds. Everything, according to the Berkeley story, is metaphysically subjective.

Berkeley's Idealism was global in that *everything* was held to be metaphysically subjective. Defenses of global mind-dependence are relatively rare these days (even in spite of the best efforts by postmodernists). In contrast, local mind-dependencies abound in the philosophical literature. For example, Saul Kripke argued, or at least asserted, that pains, if not sensations more generally, cannot exist without being experienced. If something is not felt to be painful—experienced as painful—then that something is not a pain (Kripke 1972, p. 151). On the Kripkean story, pains of which no one is aware are strictly impossible. Thus, if Kripke is right about pains, then pains are metaphysically subjective. Such a view is only a local claim of mind-dependence, that is, it is consistent with holding that other things (things besides pains) may exist independently of human awareness of their existence. Trees may fall in the woods without anyone being aware of them even though unfelt pains are impossible. States of the nervous systems of organisms may likewise be metaphysically objective: no one need be aware of them for

them to be states of nervous systems. These sorts of considerations have lead some people (including Kripke) to argue for the non-identity of pains and states of nervous systems. If pains are necessarily experienced (that is, metaphysically subjective), and states of nervous systems are not necessarily experienced (that is, metaphysically objective), then pains allegedly cannot be identical to states of nervous systems. The question of whether these sorts of arguments are sound will occupy large portions of chapter 3 (where I argue that they are not sound). I mention them here only to direct the reader's attention to uses of the notions of metaphysical objectivity and subjectivity that have relatively widespread currency in the philosophical literature.

The above examples of metaphysical objectivity and subjectivity are drawn from the philosophical literature, but the notion is not totally alien to the non-philosopher. Many readers will be familiar with the adage that "Nothing is good or bad but thinking make it so" even if they do not recall that this was said by Hamlet in Shakespeare's play. To agree with Hamlet is to hold that whether something instantiates moral properties is relative to someone's beliefs and opinions. For example, perhaps killing your uncle is wrong only insofar as someone thinks that killing your uncle is wrong. If so, then the property of being wrong is metaphysically subjective: it depends on minds for its instantiation. This is not a universally held stance on moral properties, but it is undeniably one with which many people are familiar. It is a familiar version of a claim that something is metaphysically subjective. Insofar as we understand these sorts of claims, then we also understand their denial, that is, we understand what it would be for

something to be metaphysically objective. We understand—at least in the minimal sense of getting the gist—what it means to say both that whether cats are funny is subjective and that whether cats are mammals is objective. Of course, conveying more than just a gist is the goal of this dissertation. I will have much more to say about what precisely it means to call things metaphysically objective or subjective. But in this section I have brought into view the rough outlines of the notions: notions with which almost everyone has some familiarity.

I have introduced, at least sketchily, the epistemic and metaphysical notions of objectivity. The goal of this dissertation will be to spell out a certain kind of connection between epistemic and metaphysical objectivity: I will defend the theory that what is epistemically objective depends on what is metaphysically objective. Before turning to say more about this connection, I close this section by indicating briefly what the key differences between epistemic and metaphysical objectivity are.

The epistemic and metaphysical distinctions are orthogonal. They do not simply map on to each other. A major difference between the metaphysical and epistemic senses of objectivity/subjectivity has to do with what sorts of things are either objective or subjective. The metaphysical sense of the distinction applies to everything whereas the epistemic sense of the distinction applies only to things that exhibit intentionality. Since everything that exists either depends for its existence on minds or does not, then everything is either metaphysically subjective or metaphysically objective.

Another way of seeing that the epistemic distinction is not the same as the metaphysical distinction is to see how the distinctions may cross-classify things. For example, things that are metaphysically subjective may be either epistemically objective or epistemically subjective. As mentioned above, judgments may be either epistemically objective (like maybe the judgment that the moon has no atmosphere) or epistemically subjective (like maybe the judgment that the moon is beautiful on an autumn evening). But judgments are metaphysically subjective. There are no judgments unless there are minds to do the judging, thus, in depending on minds for their existence, judgments are metaphysically subjective. It is now time to turn to the question of how things that are metaphysically subjective—things like judgments, thoughts, and beliefs—can differ in being either epistemically objective or epistemically subjective.

1.3 The Correspondence Theory of Objectivity

In discussing the idea that beauty is in the eye of the beholder I described it in terms of epistemic subjectivity. I illustrated the idea in terms of differing judgments people might form in viewing the Mona Lisa. One person may judge the painting to be beautiful and another may judge the painting to be ugly. A common view of aesthetic judgment allows that neither judgment is false, and I (for purposes of illustration, not to take a particular stand on aesthetic judgments) called both of the judgments epistemically subjective. Another way to unpack the maxim that beauty is in the eye of the beholder is in terms of metaphysical subjectivity. On this account, the painting instantiates the property of being beautiful, but its doing so depends on the reactions of certain people—people who like or otherwise have some positive attitude toward the Mona Lisa. And the Mona Lisa also instantiates the property of being ugly, but its doing so depends on the reactions of certain other people—people who (unlike the first group mentioned) dislike or otherwise have some negative attitude toward the Mona Lisa.

There are thus two ways of telling the story about the subjectivity of beauty. The epistemic story is told in terms of differing but non-contradictory judgments as to what is beautiful. The metaphysical story is told in terms of different mind-dependent properties instantiated by the Mona Lisa, properties that depend for their instantiation on different minds. The two kinds of story—the metaphysical and the epistemic—are not in

competition with each other. They are consistent with each other. Further, one of them may be explained in terms of the other, more specifically, the epistemic story maybe explained in terms of the metaphysical story. The explanation I offer is the correspondence theory of objectivity: something is epistemically objective insofar as it corresponds to something that is metaphysically objective and something is epistemically subjective insofar as it corresponds to something that is metaphysically subjective.

Of course this suggestion requires much unpacking. In particular need of unpacking are what the relation of correspondence is and what the relata—the correspondents—are. I begin by looking at the notion of correspondence as it is discussed in connection with correspondence theories of truth. The motive for comparing objectivity to truth is that there are certain suggestive similarities between the two. Just as there is a certain symmetry in objectivity, meaning that one may describe a case in either the epistemic or the metaphysical mode, a similar point is true of truth. Saying that it is true that the grass is green is one way of saying that the grass is green and we may try to understand this in terms of a correspondence theory of truth.

What follows is a toy example of a correspondence theory of truth. It is perhaps too simple by itself to be held by anyone calling themselves a correspondence theorist, but it does serve to illustrate features common to all correspondence theories. The basic components of a correspondence theory of truth are the truth-bearers (the things that are either true or false) and the correspondents (the things corresponding to the truth bearers

in virtue of which the truth-bearers are either true or false). In my toy-example, the truth bearers are sentences such as "Bill Clinton is a mammal" and "Grass is green" and the correspondents are states of affairs (construed as instantiations of properties and relations by objects) such as the state of affairs of Bill Clinton's being a mammal and the state of affairs of grass's being green. For simplicity's sake, let us consider only unquantified atomic sentences with monadic predicates—sentences of the form "*a* is *F*". A sentence of the form "*a* is *F*" is true if and only if (1) there is some individual *a* referred to by the singular term "*a*", (2) there is some property of being *F* picked out by the general term "is *F*", and (3) *a* instantiates the property *F*. The relation of correspondence obtains between the sentence "*a* is *F*" and the state of affairs of *a*'s being *F* when and only when conditions 1-3 are satisfied.

The basic picture can be extended to account for the truth of beliefs and judgments in addition to sentences by redescribing conditions 1 and 2 in terms of the singular and general concepts /*a*/ and /*F*/instead of the singular and general terms "*a*" and "*F*". The questions of what exactly concepts and terms are (e.g. Are they types or tokens? Discrete or vague?) and what the relations between them are (e.g. Are concepts just knowledge of how to apply terms? Are concepts themselves terms in a language of thought?) , are both, like much else in philosophy, matters of intense debate. Involving myself in these particular debates is something I hope to postpone indefinitely. I cannot assure the reader that nothing in the present account depends on how these debates ultimately turn out. But I do think it relatively safe to say that I presuppose nothing

agreed to be false by the majority of the debaters. I will proceed, then, by assuming the existence of both concepts and terms, and assume also that in the following account of epistemic objectivity, nothing hinges much on treating terms and concepts (and sentences and judgments) as interchangeable.

With this picture of truth as correspondence in view, I am now prepared to articulate what I take objectivity as correspondence to be. I begin by making a negative point: objectivity as correspondence is not the same as truth as correspondence. In short, the reason that this is so is because the sentence "*a* is *F*" can be objective regardless of whether it is true as already discussed above.

I will presuppose the correspondence theory of truth. This is indeed controversial. But everything in philosophy is controversial, and no philosophical work can be done without making some presuppositions. The correspondence theory of truth will be one of my presuppositions. My detractors are invited to supply an account of objectivity that does not make such presuppositions.

There are certain elements of the above correspondence theory of truth that may be pressed into the service of constructing a theory of objectivity. Just as there are truth bearers, there are objectivity bearers. Just as truth involves a relation between the bearers and the world, so will objectivity involve a relation between objectivity bearers and the world. The key difference between truth and objectivity will be what the ultimate bearers

are. The ultimate bearers of truth are sentences. This is not to deny that one does not speak only of sentences as being true, but of true theories and true stories. But this does not make them the *ultimate* bearers of truth. The truth of theories and of stories is explained in terms of the truth of sentences, not the other way around. Nothing smaller than a sentence, for instance, a singular or general term, can be a truth bearer (can be true or false). This is where objectivity departs from truth. As I shall argue, the ultimate bearers of objectivity are general terms (and concepts). This is not to deny that one also speaks of entire sentences as objective. Likewise for theories and stories. But general terms and concepts are the *ultimate* bearers of objectivity: the objectivity of sentences is explained in terms of their constituents, not the other way around.

First I want to address the question of why it is not entire sentences that are the ultimate bearers of objectivity. Consider the following, and ultimately incorrect, way of cashing out the correspondence theory. This theory defines as epistemically subjective anything that represents a metaphysically subjective state of affairs. On such an account, the sentence "*a* is *F*" is epistemically objective just in case the subject term and the general term both pick out things that are metaphysically objective.

This comports with the intuitions that the sentence "Jane is a mammal" is epistemically objective while "Jane is beautiful" is epistemically subjective. However, this version of the correspondence theory has the unintuitive consequence that the sentence "Beauty is a subjective property" is epistemically subjective. This seems

unintuitive because it violates the truism that while beauty may be in the eye of the beholder, whether beauty is in the eye of the beholder need not itself be in the eye of the beholder. Call the project of explaining why whether something is in the eye of the beholder need not itself be in the eye of the beholder "the problem of the eye of the beholder".

This problem of the eye of the beholder may be solved by treating certain sub-sentential components as the ultimate bearers of objectivity and subjectivity. Immediately, two options for the bearers present themselves: singular terms and general terms. The singular term option may be eliminated due to the following consideration. The difference between "John Smith is ugly" and "John Smith is a mammal" in virtue of which the former is epistemically subjective and the latter epistemically objective does not consist in the fact that the individual named by the subject term is metaphysically objective, since the cases do not vary in that regard. I offer, then, that the proper explication of a correspondence notion of objectivity requires only that the predicate correspond to something metaphysically objective. Metaphysically objective objects need not be referred to in singular sentences. Instead, I offer, what makes a singular sentence of the form "*a* is *F*" epistemically objective is that the property *F* named by the predicate term "*F*" is metaphysically objective. Thus I call the correspondence theory of epistemic objectivity that I advocate "the predicational theory of epistemic objectivity". Below I sketch how the predicational theory deals with singular and quantified sentences and the corresponding propositional attitudes.

To keep the discussion at a manageable length, I will define epistemic objectivity only for atomic sentences and their corresponding propositional attitudes. I will also restrict my attention to sentences containing binary or monadic predicates and attitudes employing binary or monadic concepts.

The inscribed or uttered sentence token "*a* is *F*" is epistemically objective just in case the property *F* is metaphysically objective. "*a* is *F*" is epistemically subjective just in case the property *F* is metaphysically subjective. Whether "*a* is *F*" is epistemically objective is independent of whether *a* is metaphysically objective. Quantified sentences of the form "$(\exists x)$(*x* is *F*)" and "(x)(*x* is *F*)" are epistemically objective just in case the instantiation of the property *F* is metaphysically objective. The objects quantified over need not themselves be metaphysically objective.

I extend the treatment of the epistemic objectivity of sentences to apply to propositional attitudes such as belief. Propositional attitudes, I will suppose, require the possession and deployment of concepts corresponding to the terms used in the sentences that are their expression. My believing that *a* is *F* involves my possessing and deploying the subject concept of *a* and the predicative concept of being *F*. The same predicative concept is employed in my quantificational beliefs that all *x*'s are *F* and that some *x*'s are *F*. The epistemic objectivity of propositional attitudes such as belief that *P* and the

perception that *P* is due to my possessing and deploying a predicative concept that picks out a metaphysically objective property.

I want to take some time here to briefly remark upon the superiority of the correspondence theory over its competitors: the indexical and consensus theories. Much of the dissertation is taken up with my case against competitors, but I want to make some brief foreshadowing remarks here.

The first issue to consider is how (and whether) the three theories account for the subjectivity of "cats are funny". Both consensus and correspondence theories do well here, but the indexical theories immediately run into trouble. "Cats are funny" is third personal, not first personal. It is devoid of any obviously indexical/demonstrative terms such as "I" or "this".

Now, the indexical theorist may respond that perhaps "funny" is itself indexical. It may be viewed as indexical insofar as it picks out different things when used by different people: it may mean something quite different in my mouth than in yours. When I say "cats are funny" that is equivalent to "cats amuse me". When you say "cats are funny" you don't mean that cats amuse Pete Mandik, but instead that cats amuse you. Thus insofar as "funny" has meaning determined in part by context, it is indexical. The problem with this kind of response is that it can be used to show that just about any representational state is indexical, and thus subjective, since for any term there is a sense

in which its meaning depends on context. "Cats are mammals" turns out to be indexical insofar as the meaning of "mammal" is indexed to Engish speakers in the actual world. It is possible that there is another language (or an alternate English) in which "mammal" means what we mean by "vegetable". Allowing the notion of indexicality to be so widely applicable is bad for the indexical theorist considering their primary motivation for adopting the indexical theory, namely, constructing a physicalist response to the knowledge argument. Indexical theorists want to allow that there is a meaningful if not crisp objective/subjective distinction whereby statements in natural sciences such as physics turn out to be objective. But detecting subjectivity wherever there is context sensitive meaning is antithetical to such a project. As I will discuss at further length in Chapter 3, much of the superiority of the correspondence theory of the indexical theory hinges on issues surrounding the knowledge argument and the subjectivity of consciousness. Similar considerations also count against the consensus theory, since it is absolutely silent on such issues.

1.4 The Paradoxes of Subjectivity

There is a tradition in philosophy of giving the following partial answer to the "what sorts of things are objective/subjective" question of application: properties. On this tradition, secondary-properties are subjective properties, primary properties objective properties. A different answer to the "what sorts" question says that conscious

experiences are subjective, material bodies are objective. In both traditions, there is a sub-tradition of formulating accounts of subjective things in terms of accounts of objective things. In both traditions, an informal paradox arises, I call it "the paradox of subjectivity". The crux of the paradox begins with a definition of subjective properties in terms of objective properties and ends with a conclusion that something is both objective and subjective.

In the case of secondary qualities, the seeds of paradox are sown as follows. Primary qualities of objects include the way objects occupy and move through space: the shapes and motions of objects. Any qualities definable in terms of the occupation of and movement through space were also considered primary, and thus objective. Thus velocities, accelerations, attractive and repulsive forces all fall under the heading of "primary". What room, then, is there for the secondary qualities? The classical conception of secondary qualities is as of the powers in objects to cause certain reactions in observers. But powers exercised by one body on another are as primary—and thus as objective—as shape and motion. Taking color as a paradigm of secondary qualities, the informal paradox may be formulated as follows:

> Premise 1. Barney's being purple is subjective
> Premise 2. Barney's being purple is just Barney's having a disposition to cause certain a reaction in observers.
> Premise 3. But whether Barney has a disposition to cause certain observers to represent him as purple is entirely objective
> Conclusion: Barney's being purple is objective

Note what reference to observers may conceal or assume, namely, that being an observer—a subject of sensory experience—is something reducible to primary properties. Perhaps what separates powers to cause reactions in observers from powers to cause reactions in other entities is that observers have conscious experiences. Perhaps what is primal to the subjectivity of secondary qualities is not merely that they are dispositional, but that they are dispositions concerning the effects on conscious experience. Perhaps conscious experience is the prototype of subjectivity and secondary qualities are subjective only derivatively. It is still open to such a view to hold that experiential properties are reducible to physical properties. Mention of the physicalization of experience leads to the second version of the paradox of subjectivity.

> Premise 1. What it's like to be a bat is subjective
> Premise 2. What it's like to be a bat is just the bat's being in a particular brain state
> Premise 3. Whether the bat is in a particular brain state is entirely objective
> Conclusion: What it's like to be a bat is objective

Both paradoxes may be dissolved by the predicational correspondence theory of objectivity. I begin by discussing the paradox as it arises for secondary qualities. The key insight for the solution of the paradox of secondary qualities is the following. Beauty may be in the eye of the beholder, but whether beauty is in the eye of the beholder need not itself be in the eye of the beholder. Whether spinach is yucky is up to us, but whether it is up to us need not itself be up to us. Echoing the initial formulation of the paradox in

terms of the supposition that anchovy ice cream instantiates the secondary property of being disgusting yields the following.

> Premise 1. Anchovy ice cream's being disgusting is subjective
> Premise 2. Anchovy ice cream's being disgusting is just Anchovy ice-cream's having a disposition to cause certain a reaction in observers.
> Premise 3. But whether Anchovy ice cream has a disposition to cause certain observers to represent it as disgusting is entirely objective
> Conclusion: Anchovy ice cream's being disgusting is objective

The dissolution of the Paradox for Secondary Qualities involves restating the key points about objectivity and anchovies as follows.

1. " Anchovy ice cream is disgusting" is epistemically subjective
2. " 'Anchovy ice cream is disgusting ' is epistemically subjective" is epistemically objective
3. Disgusting-ness is metaphysically subjective
4. The property-of-being-metaphysically-subjective is metaphysically objective

1-4 are consistent. This consistency is enabled by both (i) distinguishing between epistemic and metaphysical objectivity/subjectivity and (ii) recognizing that properties and predicates are themselves bearers of metaphysical and epistemic objectivity, respectively. These remarks on how to dissolve the paradox of subjectivity support the advantage of treating the ultimate bearers of epistemic objectivity as predicates and metaphysical objectivity as properties. I postpone until chapter 3 discussion of the paradox as it relates to the subjectivity of conscious experience.

In the remainder of this section I discuss a possible objection to the account given so far. [1] Brains and brain states, many suppose, count among the objective things. Some brain states may also be mental representations, on many physicalistic accounts of mind. Mental representations, let us assume, are individuated by their representational contents. Assume further that some representations are about themselves: "This sentence has six words in it" though false is about itself. Magritte's painting "Ce n'est pas une pipe" is about itself. A carpet sample is about, among other things, itself (Goodman 1976, pp. 57-59). The dictionary definition of "definition" is about, among other things, itself. Mental representations can be about themselves too. If you sincerely attempt to form an answer to the question "What are you thinking about?" then you will have tokened a mental representation about itself, since you have thought a thought about itself. If you wonder to yourself how many neurons are in the supervenience base of your occurrent mental state, then you again token a mental representation about itself.

Mental representations that are about themselves are metaphysically subjective. They exist only if they are represented and are thus mind-dependent. Brain-states that are about themselves are likewise mind dependent. But, as mentioned earlier, brain states are objective. One and the same thing seems to be metaphysically objective and metaphysically subjective. Is this contradictory?

---

[1] The crux of the objection is due to Jesse Prinz, personal communication.

The apparent contradiction can be shown to be illusory by appealing to a distinction between superordinate and subordinate properties (like being colored vs. being red).The property of having a brain in a particular state is metaphysically objective, since organisms can have brain states without being represented as having brain states. The property of having a self-referential brain state, is, of course, metaphysically subjective. The intuition that the property in question is metaphysically objective is generated by shifting attention from it to the superordinate that it falls under. The property of having a self-referential brain state seems, then, not to be an example of something that is both metaphysically objective and metaphysically subjective. Some brain states turn out to be metaphysically subjective. This is quite compatible with the fact that the property of having a brain state is metaphysically objective.

Is this a retreat to property dualism? No. I have made an appeal to different properties. But it's an appeal that's consistent with physicalistic property monism. Being an electron and being a subatomic particle are sub- and super- ordinates, respectively. But both are physical properties. Likewise then, for being felt as pain and being a brain state.

1.5 Things to Come

In this chapter I have initiated an inquiry into the ways in which representational states may be objective or subjective, more specifically, how they may be either epistemically objective or epistemically subjective. I unpacked and began arguing for an account of epistemic objectivity: the correspondence theory whereby epistemically objective things represent metaphysically objective things. I argued for a particular version of the correspondence theory whereby predicational representations are the ultimate bearers of epistemic objectivity and subjectivity. Much of my argument in favor of the predicational theory centered on the way in which it is able to dispel certain paradoxes of subjectivity. Further support for the theory will be marshaled in subsequent chapters. In the next chapter I discuss the objectivity of representation against a backdrop of naturalized accounts of representation.

Chapter 2: Egocentric and Allocentric Mental Representations

2.0 Introduction

In the previous chapter I began arguing for the view that epistemic objectivity arises when a representation represents something metaphysically objective. I want now to spell out this account within the context of the naturalization of representation. Some caution needs to be taken, however, with what exactly the naturalization of representation amounts to. For many philosophers, the project of naturalizing representation involves attempting to analyze the concept of representation into necessary and sufficient conditions stateable using only the vocabulary of the natural sciences (paradigmatically, physics). This is not the notion of the naturalization of representation that I am particularly interested in. In section 1.1. I eschewed analysis in favor of the Quinean notion of explication. Another Quinean notion that I am fond of is the conception of philosophical naturalism not as a special brand of conceptual analysis, but instead as the view of the relation between philosophy and the natural sciences as one of continuity. Naturalizing representation, then, will involve leaning heavily on scientific uses of the notion of representation, especially as it is used in cognitive science and neuroscience. My allegiance to this vision of naturalism does not mean that I hold no regard for the

results of the conceptual analysts: indeed, many of their insights are worthwhile and workable, even if their larger program is not one I have a lasting interest in.

A version of the epistemic objective/subjective distinction appears in explanations of cognition under the guise of a distinction between allocentric and egocentric representations. Several philosophers and psychologists interested in the topic of spatial representation have identified a distinction between egocentric and allocentric representations of space and described this distinction as one between subjective and objective ways of representing space (see, for example, Brewer 1996; Campbell 1996; Evans 1982, O'Keefe 1996). In this chapter I flesh out this suggestion along the lines of the correspondence theory. I begin by examining the notions of perspective and point of view that will be crucial in the discussion in Chapter 3 concerning the subjectivity of conscious experience. In this chapter I supply an explication of egocentric and allocentric representations. The crux of the explication involves understanding a distinction between, on the one hand, representations that are perspectival or embody a point of view, and one the other hand, representations that abstract away from perspectives or points of view.

2.1 Mental Representation and Pictorial Perspective

What does it mean for a representation to have or be from a point of view? A philosophically popular kind of answer to the above question begins by latching on to the literary usage of the phrases "point of view" and "perspective" (See, for example, Lycan 1996). Prose written in the so-called first-person point of view employs indexical terms to refer to the author or alleged author. One famous example would be the opening line of Melville's Moby Dick: "Call me Ishmael". The second person point of view uses indexical/demonstrative terms to address the reader, as in the sentence "You may be wondering what will happen next". Typical instances of prose written in the third person point of view are devoid of indexical and demonstrative terms, as in "On November 19, 1999 Pete Mandik purchased a portrait of a platypus".

There is a less literary and more literal way of understanding perspective and point of view. This sense also has more promise for neuroscientific research and for dealing with the knowledge argument, as will be discussed in chapter 3. This other way representations may be from a point of view would be in the rather literal sense that pictorial representations embody a point of view. To see how pictorial representations embody a point of view consider two photographs of the same object taken from two different angles. Compare, for example, two photographs taken of a person's face: The first may be head-on, the other may show the head in profile. The camera that produced

the photos occupied two different points of view with respect to the person's head. Let us, for convenience, describe possible locations of the camera around a horizontal plane in terms of clock positions and distances from the center of the person's head. So, the full frontal photo would be from the twelve o clock position at a distance of 6 feet, where the profile shot would be from the 3 o'clock position at a distance of 6 feet (see figure 1).

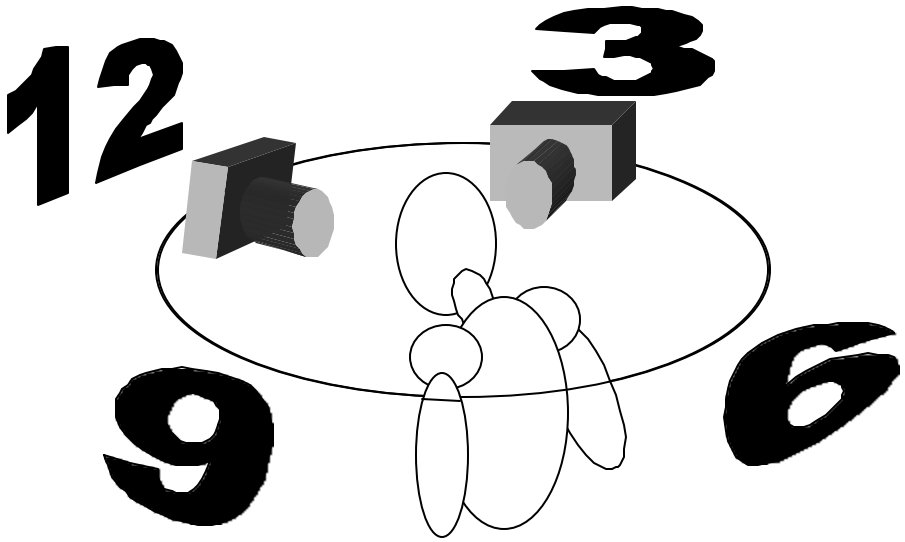

Figure 1.
Cameras occupy literal points of view with respect to their subjects.

The different positions and orientations of the camera specify the points of view of the camera with respect to the subject. The representational contents of the photographs produced include content about these points of view. This is why we can tell, for instance, by looking at the photographs of face profiles, whether the camera was at the three o'clock or nine o'clock positions. See figure 2.

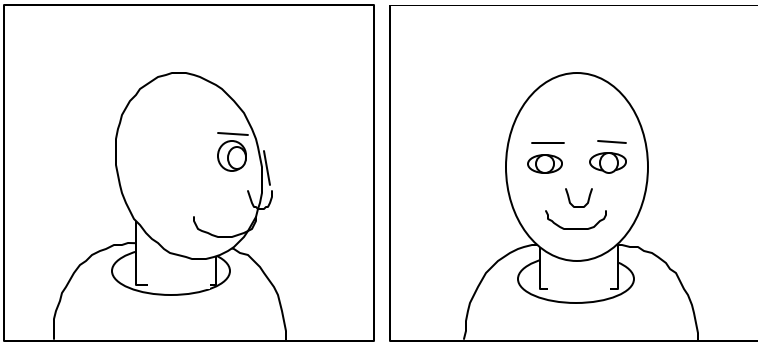

Figure 2.
Pictures have as part their representational content relations that the viewer would bear to the subject. This is why we can tell, e.g. the position of a camera with respect to the subject just by looking at the photographs.

Note that the pictorial sense of "point of view" is rather literal but not totally literal. Few imagistic representations contain enough information to specify a particular point of view. For example, a typical map of Chicago presents a "bird's eye view" of the city, but abstracts away from any information that would specify a point of view positioned over the Sears Tower as opposed to the Hancock Building. This is not to say that the map abstracts entirely away from point of view: it does contain enough information to specify that one is viewing the city from above it rather than below it and that one is viewing the city from a location over a part of the United States as opposed to a location over Japan. For another example of how images may embody a point of view without specifying a literal point of view, consider drawings in oblique perspective as opposed to drawings in vanishing-point perspective (see figure 3).

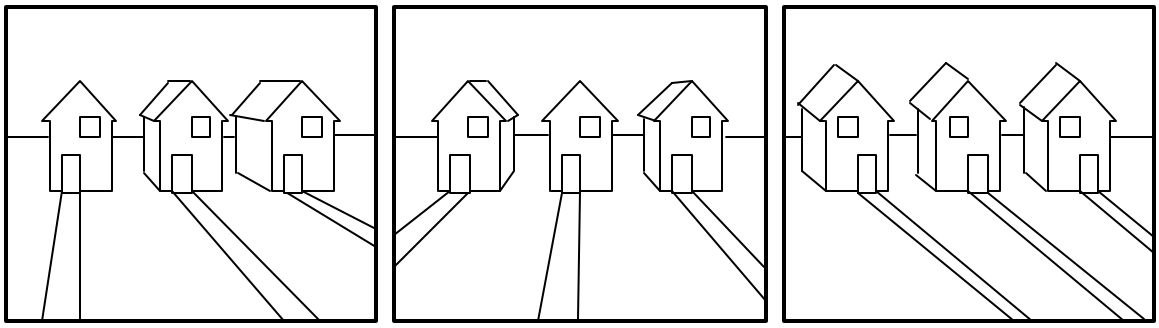

figure 3a figure 3b figure 3c

Figure 3. Figures 3a and 3b are drawn in vanishing-point perspective, thus positioning the viewer in front of the first and second houses, respectively. Figure 3c is drawn in oblique perspective and thus does not specify a point of view in front of any particular house.

Drawings in vanishing point perspective represent parallel lines that are not perpendicular to the viewer's line of sight as converging on one or more points (the vanishing-points). Figures 3a and 3b are drawn in vanishing-point perspective. The vanishing-point in figure 3a is located at the peak of the first house from your left, and in figure 3b, at the peak of the central house. In oblique perspective, parallel lines not perpendicular to the viewer's line of sight are parallel in the drawing, but drawn at non-right angles to the lines that are perpendicular to the viewer's line of sight. Figure 3c is drawn employing oblique perspective. Drawings employing oblique perspective abstract away from the position of the viewer more than do pictures employing vanishing point perspective. Figure 3a specifies a point of view in front of the first house and figure 3b specifies a point of view in front of the second house. Figure 3c, drawn in oblique perspective, abstracts away from which of the three houses the viewer is in front of, since, unlike in figures 3a and 3b, the sides of the three houses occupy equal portions of the picture plane.

Nonetheless, Figure 3c does not abstract away entirely from point of view as is evident from the fact that we can see only the front, top, and one side of each house.

The pictorial sense of perspective can be understood as a not necessarily implicating a precise point of view. The pictorial notion of perspective is thus just the notion that part of the representational content of the picture includes relations between what is pictured and the viewer. Such relations would include being in front of or to the side of a face or a house. Other relations would include being closer to one pictured object than another. Another way of seeing how pictorial perspective abstracts away from literal, geometric, points of view is in the notion of what is known as aerial or atmospheric perspective. Pictures employing atmospheric perspective depict things in the far distance as being fainter, hazier, and bluer than things that are closer to the viewer. The device of atmospheric perspective exploits facts about the behavior of light in the atmosphere to depict relations of distance between the viewer and what is pictured, and is a frequently employed device in landscape painting. The sort of relations captured by atmospheric perspective abstract away from the precise location of the viewer. For example, it would be difficult to ascertain from atmospheric perspective which of the houses in figure 3a the viewer was in front of.

Pictorial perspective is not only a property of photographs, drawings, and paintings. Mental representations also exhibit pictorial perspective. In describing pictures as involving pictorial perspective I said that the representational contents of the pictures

include relations between the things pictured and the viewer. Extending this account of perspective to mental representations yields the thesis that some mental representations include in their contents relations between the representing subject and that which is represented. Thus such representations are egocentric (self-centered) since they represent relations that things bear to the representer. I turn now to the case that there are such mental representations.

The point that mental representations exhibit pictorial perspective can be bolstered by both phenomenological and empirical scientific considerations. Consider first the phenomenological support. Just as two different photographs of the same house may be from two different points of view, so may two different visual experiences be from two different points of view. What it is like to look at the house from a point to the north of it may be quite different from what it is like to look at the house from a point to the west of it. Thus the perceptual representations involved in the two different experiences exhibit pictorial perspective. We need not actually be perceiving a house in order to have mental representations that exhibit pictorial perspective. One may instead dream of seeing a house from one point and then another and notice the differences in point of view. One may also imagine looking at the house from the north and then the west and in each instance a difference in point of view is introspectible in imagination. These points based on introspection may be enhanced by evidence from psychological studies, to which I now turn.

Much psychological research in the past several decades has concerned the nature of mental images and speaks to the issue of the existence of perspectival representations. A classic example is due to R. N. Shepard and his colleagues (Shepard & Cooper, 1982). These researchers had subjects look at simultaneously presented pairs of objects. The second member of each pair was either the same as the first or a mirror image. Further, pair members could differ from each other in their rotations in depth and in the picture plane. The researchers found that the time it took for subjects to make "same or different" judgments increased monotonically with increases of rotational displacement between pair members. Shepard et al. took this reaction time data as evidence that subjects were rotating mental images to see if they would match the stimulus. The evidence that Shepard et al collected also serves as evidence for the existence of pictorially perspectival mental representations. A mental image at any given stage of a rotation constitutes a perspectival representation because at each point in rotation, the image represents what the object would look like from a particular point of view.

Some theorists have postulated that mechanisms similar to those postulated for image rotation may be at work in visual object recognition. Humans recognize visually presented three-dimensional objects with only two-dimensional projections on the retina as a guide. Somehow, we are able to recognize objects seen from unfamiliar viewpoints, that is, based on unfamiliar projections onto our retinas. Certain studies of the accuracy and reaction times in visual recognition tasks implicate perspectival representations. Such studies typically examine the reaction times and accuracy of recognition judgments

of objects seen from unfamiliar viewpoints. In such studies, average length of reaction time and judgment accuracy varies monotonically with the degree of rotational deviation (in depth or on the picture plane) from familiar views of an object. These correlations are taken as evidence for the hypothesis that visual object recognition is mediated by a normalization mechanism. The stored representation of an object is one or more encoded 'views' that encode only two-dimensional information based on previous retinal projections. Recognition of familiar objects seen from unfamiliar viewpoints involves a match between a stored view and the perceptual view via a normalization mechanism which compares the views (e.g., Bülthoff & Edelman, 1992; Shepard & Cooper, 1982; Ullman, 1989). For example, this might involve mentally rotating an image (Shepard & Cooper, 1982). Object recognition, as well as imagery, may involve perspectival representations.

Perspectival representations also surface in accounts of navigation. Just as there is evidence that an object may be recognized better from one point of view than another, so may a destination be better arrived at from one starting point than another. One rich body of research on navigation involves lesion studies of rats. One kind of experiment in rats concerns the performance of lesioned rats in the Morris water maze. The Morris water maze is an apparatus filled with water in which rats can swim. Objects such as small platforms can be placed in this arena. Milk powder can be added to the water to make it opaque, and the level of the water can be adjusted so that when a platform is submerged it is not visible to rats swimming in the maze. In Eichenbaum et al. (1990) a

water maze was set up such that rats had to swim to a platform visible during training trials, but occluded by the opaque water in the testing trials. Varied visual stimuli were positioned around the maze to serve as orientation cues. The experimenters trained intact and hippocampal-system damaged rats to swim to the platform from a given start location. During test trials, both the intact and damaged rats were able to swim to the platform if they were started from the same location as in the learning trials. However, the performances of the intact and damaged rats diverged widely when they were started from novel locations in the water maze. During test trials, intact rats were able to navigate to the platform from novel start locations, whereas the hippocampal damaged rats required much longer to find the platform, and sometimes never found it during the test trial.

Much discussion of these sorts of investigations of the neural bases of rat navigational abilities has concerned the proposal that the hippocampus is a locus of allocentric ("other centered") representations of the spatial locations in the maze and without which lesioned rats must rely on merely egocentric ("self centered") representations of spatial locations (see e.g. Okeefe and Nadel 1978). Evidence of egocentric and allocentric representations of spatial locations is not confined to studies of rats. For example, Feigenbaum and Rolls (1991) recorded the electrical activity of individual neurons in the hippocampus of macaque monkeys. In their study, they looked for neurons that were maximally responsive to the spatial location of a visual stimulus. They then changed the spatial relation of the monkey with respect to the stimulus so that,

although the stimulus had not moved, it projected to a different part of the monkey's retina. Cases in which activity in neurons was still maximally responsive to stimuli in that location regardless of what part of the retina the stimulus projected to were regarded as allocentric representations of that spatial location. In contrast, neural activity maximally responsive to a spatial location defined relative to the site of retinal projection is regarded as an egocentric representation of that location. Feigenbaum and Rolls (1991) report that the majority (but not all) of the cells in hippocampus that they investigated were allocentric representations of spatial location. Where hippocampus has been widely implicated as the locus of allocentric representations of spatial locations, many kinds of egocentric representations of space (including, for instance, head-centered and shoulder-centered representations) have been localized in regions of the posterior parietal cortex.[2] However, questions of where egocentric representations are are not as interesting as questions of what egocentric representations are. I propose to explore such questions against the backdrop of naturalistic theories of representational content.

Developing and defending theories of content, especially ones consistent with a naturalistic or physicalistic ontology, is a central topic in current philosophy of mind. A theory which seeks to explain what it is for a representational state to be about something—to provide an account of how states and events can have representational

---

[2] See Stein 1992 and Milner and Goodale 1995 for general discussions of egocentric representations in parietal cortex.

contents—is known as a psychosemantics. A physicalist psychosemantics seeks to do this using resources of the physical sciences exclusively. What precisely count as the physical or natural sciences is difficult to specify without becoming embroiled in controversy, but suffice it to say that the neurosciences count among them.

One approach to psychosemantics is informational psychosemantics. It invokes the causal relations obtaining between the state and the object it represents to ascribe content. A physical state represents birds, for example, just in case appropriate causal relations obtain between it and birds. At the heart of informational semantics is a causal account of information (Dretske, 1995). Red spots on a face carry the information that one has measles because the red spots are caused by the measles virus. A common criticism of informational semantics holds that mere causal covariation is insufficient for representation, since information (in the causal sense) is by definition always veridical while representations can misrepresent. A popular solution to this challenge invokes a teleological notion of function. A brain state represents *X* by virtue of having the function of carrying information about (being caused by) *X* (Dretske 1988). The teleological notion of function allows for a distinction between having a function and performing a function. For example, a person's eyes may have the function of seeing even though a particular person is blind and thus has eyes that fail to perform what is nonetheless their function. The crux of the teleological/informational solution to the problem of misrepresentation is to analyze having representational content as having the function of

carrying information, a function that may be had by something even in cases in which it fails to perform its function, and thus, misrepresents.

While informational psychosemantics offers these conditions as both necessary and sufficient for representations, here I need only to regard them as sufficient. Below, in section 2.2 I argue against their necessity and offer an additional set of sufficient conditions: *X* may represent *Y* if *X* has the function of causing *Y*. As I will discuss in 2.2 this (procedural instead of informational) account is most appropriate to accommodating the view, widespread in neuroscience, that activation of areas in motor cortex represent efferent events. We have here then two kinds of condition sufficient for *X* to represent *Y*: one expressed in terms of *X*'s causing *Y*, one in terms of *Y*'s causing *X*. I shall want to combine the two and say that it is sufficient for *X* to represent *Y* that *X* has the function of being causally related to *Y* (alternately: causally covarying with Y), thus covering both the informational (affector) and procedural (effector) cases. I make no claims on these sufficient conditions also being necessary. I have neither the resources nor the need here to offer a full-blown account of representation, an account that can supply conditions for all kinds of mental representations. My main concern here is with accounts adequate for representation as it figures in sensory experience. I will assume without further argument that causal covariational accounts are adequate.

The neurobiological paradigm for causal covariational semantics is the feature detector: one or more neurons that are (i) maximally responsive to a particular type of

stimulus, and (ii) have the function of indicating the presence of a stimulus of that type. Examples of such stimulus-types for visual feature detectors include high-contrast edges, motion direction, and colors. A favorite feature detector among philosophers is the alleged fly detector in the frog. Lettvin et al. (1959) identified cells in the frog retina that responded maximally to small shapes moving across the visual field. The inference that such cells have the function of detecting flies and not just any small moving thing is based on certain assumptions about the diet and environment of frogs, thus satisfying (ii) as well as (i). Using experimental techniques ranging from single-cell recording to sophisticated functional imaging, neuroscientists have recently discovered a host of neurons that are maximally responsive to a variety of stimuli. Among these are neurons that are particularly sensitive to the spatial locations of stimuli. In some cases the neurons are responsive to locations relative to the subject, thus giving rise to perspectival, or egocentric, representations of spatial locations. In other cases, the neurons are responsive to locations independent of the relations between the location and the subject, thus giving rise to non-perspectival, or allocentric, representations of spatial locations.

Now we are in a position to see how the notion of perspectival representations such as egocentric representations of locations may be accommodated by a casual covariational psychosemantics. A subject *S* has a perspectival representation *R* of *X* if and only if the representational content of *R* includes relations *S* bears to *X*. Cashing out the notion of representation in terms of the teleological and causal covariational account described above yields the following formulation. A subject *S* has a perspectival

representation *R* of *X* if (but maybe not only if) *R* has the function of causally covarying with *X* and relations *Z*1-*Z*n *S* bears to *X*. In the case of spatial representations, on which I will focus for now, the relations in question will be spatial relations. Later I will generalize this definition of perspectival representation to non-spatial sensory modalities. One class of spatial perspectival representations is provided by neurons with retinocentric receptive fields. Such a neuron, whether in cortex or in the retina itself, demonstrates a pattern of activity maximally responsive to the occurrence of a specific kind of electromagnetic radiation in a certain spatial location defined relative to the retina. It is a plausible and widespread assumption that activity in neurons with retinocentric receptive fields represent (or "encode" or "code for") luminance increments in retina-relative spatial locations. If this assumption is correct, then we can see how such neural representations conform to the account of perspectival representations. In this example *R* is a certain kind of activity in a certain neuron in *S*'s nervous system, *X* is a luminance increment and *Z*1-*Z*n include the spatial relations *X* bears to *S* (especially spatial relations to *S*'s retina). For another example, consider neural activity that represents goal locations for saccades. Analogous to the receptive fields of sensory neurons, motor neurons have what we may call "effective" fields. The effective field of a neuron may be a region in space that an organism may move or reach toward in response to activity in a particular neuron. There are neurons that control saccades that have as effective fields head-relative spatial locations. If it is correct to speak of activity in such neurons as "representing", "encoding" or "coding for" head relative spatial locations, then we have another instance of perspectival representations. Activity in such motor neurons

represents the movement of the eye toward a location in space defined relative to the subject's head. Such neural activations do not simply causally covary with the movement of the eyes to a certain location, but the movement of eyes to a certain location defined relative to the subject, and in this instance, relative to the subject's head. And if these activations have the function of causally covarying with these subject-relative locations then they constitute perspectival representations of spatial locations.

A brief word needs to be said about the compatibility of mental imagery and the causal covariational account of representation. Some researchers favor an account of imagery whereby images represent in virtue of resembling that which is represented (See Kosslyn 1994). There is much literature on this issue, and suffice it to say, few agree that resemblance is necessary for representation, even in cases where the representations are images. To illustrate the point, consider finding a creature inside of which we found something that looked like this:_/\/_. Suppose that we wondered whether this constituted a representation of something. It resembles a mountain range. Might it be a representation of a mountain range? It also represents saw-teeth, a row of evergreen trees, and abandoning visual resemblance, we may say that it resembles a noise with a certain waveform. Which does it represent? Resemblance, as many have pointed out, underdetermines representation, and even in cases in which representations do resemble what they represent, functional causal covariation may be called in to do the job of disambiguation. For instance, a photograph of Joe equally resembles Joe and Joe's

identical twin Moe. But the photograph is a photograph of Joe and not Moe in virtue of Joe's position in the causal chain that led up to the creation of that photograph. These points about the role of causal covariation in determining the contents of imagistic representations are consistent with a common analysis of mental imagery. According to this analysis, imagery is the off-line utilization of perceptual (and perhaps motor) processes that are typically used on-line (see for example, Grush 1997). Since that is all there is to imagery, imagery need not resemble what it is an image of. Such a view allows for imagery in non-visual modalities such as olfaction, where the possibility of resemblance between the representation and the represented seems obscure. I mention these points not to settle any ongoing controversies regarding mental imagery, but only to show that the existence of imagistic mental representations is neither necessarily nor obviously incompatible with accounts of representation that make having the function of causally covarying with *X* sufficient for representing *X*.

Another alleged incompatibility of imagery and causal psychosemantics can be shown to be merely apparent. Recall that part of the account on offer is that a subject *S*'s representation *R* of *X* is perspectival if it has the function of carrying information about relations between *S* and *X*. Rick Grush (personal communication) objects that this kind of account is inadequate for the perspective embodied in imagery on the grounds that

> Taking S to be the actual real S then entails that my imagining seeing the Eiffel Tower right now from the north is not perspectival, because it caries no information about the relations between S (the actual me sitting here in the [. . .] café) and the Eiffel Tower.

Grush may be correct that the event of imagining does not carry information about the relations between S and the Eiffel Tower. On a particular occasion one may be caused to imagine the Eiffel Tower by something other than the Eiffel Tower, and thus that particular imagining would not carry information about the Eiffel Tower. However, on the account of imagery sketched in the previous paragraph, the event of imagining involves a state that has the function of carrying information about the relations between S and the Eiffel Tower. Imagining involves running off-line what is run on-line during perception: states that are supposed to carry information in the perceptual case may also be pressed into service for off-line imaginings. Thus the off-line states employed in imagining the Eiffel Tower owe their representational content to the information they are supposed to carry in the on-line perceptual case.

I call the analysis of perspective I offer "pictorial" in order to contrast it with the analysis of perspective offered by Lycan (1996) and others in terms of indexicals and the literary convention of first-person point of view. Pictures are the prototypical instances of representations with pictorial perspective, but it is important to emphasize that they are not the only instances. Pictures of a car from two different points of view may be encoded as bitmaps, which may themselves be translated into strings of ones and zeros or sentences describing the occupants of every cell in the bitmap's two-dimensional array. One may find it natural to suppose that such resulting strings of numerals and sentences may retain the representational content of the pictures from whence they came without themselves being pictures. This is not to say, however, that bitmaps have all and only the

content of the images they encode. There may be some differences in the representational contents of the images and the bitmaps. However, despite possible differences of representational content, there are also considerable similarities. If a bitmapped photograph represents a car, then the corresponding numerical string does too. After all, the picture of a car is recoverable from the bit string. And if the bitstring retains the representational content of being about a car, then there seems no reason to deny that the bitstring also retains the representational content of being about a car as seen from a particular point of view. Thus, if a picture is perspectival , then its corresponding bitstrings are perspectival, even though the corresponding bitstrings are not pictures.

Bitstrings are not the only instances of representations with pictorial perspective that are not pictures. Activations in neurons with retinocentric receptive fields are another example. Such neural activations (are thought to) represent the occurrence of stimuli at spatial locations defined relative to the retina. I take it as obvious that the activation of a single neuron is not a picture even in cases in which such activation may be a spatial representation. Below I will discuss the possibility of perspectival representations of temperature, thus giving yet another example of representations with pictorial perspective that are not pictures.

Representations with pictorial perspective are not necessarily pictures. Nor are they necessarily indexical. Above I mentioned that I intend the pictorial analysis to contrast with indexical analysis, but have not shown that this is indeed the case. The reason why representations with pictorial perspective are not necessarily indexical has to do with particularity. Indexicals necessarily involve particularity in a way that non-indexicals and egocentric representations do not. The representational content of the utterance "I am here now" picks out a particular individual at particular location at a particular time. Even when indexicals and demonstratives are used to pick out universals, as in saying "this shade of red" while holding up a chip of paint, reference to the universal (the shade of red) piggy-backs on the particulars that secure the indexical content: the particular paint chip held by the particular individual at a particular time. In contrast, a picture can exhibit perspectival content without picking out any particular. A drawing of the 1991 Chevy Cavalier may be used not to represent any particular 1991 Chevy Cavalier. It may instead be used to represent a corresponding universal, say, the general category that all and only 1991 Chevy Cavaliers belong to. Nonetheless, the picture, in being a two-dimensional representation of a three-dimensional object, represents the car from one or another point of view. For instance, the picture may show the front of the car but not the back. Of course, the conventions of photography may have an indexical element: a photo represents me and not my twin in virtue of being appropriately caused by me, not my twin. But a drawing of the general body style of the 1991 Chevy Cavalier need not have its representational content determined in the way a photo of a particular 1991 Chevy Cavalier would. But both would be perspectival—both

would implicate a point of view. Thus, there is a kind of perspective—pictorial perspective—that is underdetermined by indexical content. Something may exhibit pictorial perspective without being or containing an indexical.

2.2 Motor Control and Spatial Representation

I next turn to a possible objection to the explication of egocentric spatial representation of space set forth above, namely that so-called egocentric representations of space are not representations at all, but instead mere motor commands. I argue for a construal of the role of motor control in spatial representation whereby motor commands can be, or at least figure in, representations of space.

The objection that so-called egocentric representations are mere motor responses and thus non-representational derives most of its force from contemplating the role that such representations are alleged to play in rat navigation. Rats thought to be lacking allocentric representations, such as those in the maze learning study of Eichenbaum et al (1990) described above, exhibit a minimal memory of previous trials. But they are able to find their way to goal locations only be retracing their steps, so to speak. This leads to the suggestion that maybe what the rats are representing isn't a goal location, but merely remembering a previous pattern of motor responses to a given stimulus. Indeed, even those who are disinclined to deflate the notion of egocentric representation in favor of

motor response are still inclined to see motor response as a constitutive aspect of egocentric representation (e.g. Milner and Goodale, 1995 and Kirsh 1996). My aim in the present section is to spell out how egocentric representation may be closely related to motor response without caving into the objection that so-called egocentric representations are *mere* (non-representational) responses.

As mentioned above, typical causal accounts of representation rely on the notion of information. I want to contrast the informational story with what I shall call a procedural story: a story whereby representations can be about their effects instead of their causes. I will say more about the details of the procedural story in the next section. For now let me just flag the following. I think that the informational story is plausible for some mental representations—I just do not think that it will be the truth about all representations. I think that some representations will best be handled by a representational story that has procedural as well as informational components. I focus on mental representations as they figure in our perceptual experience of spatial properties.

Our senses of taste and smell seem not to have much to do with the spatial properties and relations of physical objects. In marked contrast, our senses of hearing, touch, and—most of all—sight, present us with richly structured spatial manifolds. We hear sounds as coming from various directions in relation to our bodies. And most notably and importantly for fans of stereophonic recordings of music, we can hear sounds as coming from either the right or left sides of the spaces our bodies occupy. Our senses

of touch and sight allow us to discriminate the shapes of objects, as well as detect the various distances between each of the objects.

One particularly noteworthy feature of vision is that it functions as a "distance" sense: the properties sensed via that modality are presented as being located at various distances from the surfaces of our bodies, even though the sites of transduction are located in our retinas. In contrast to our retinas, our tongues present their stimuli (flavors, temperatures, and textures) as occurring right at the site of transduction. Interestingly, our sense of touch, while typically not a distance sense, can sometimes function to the contrary. The most familiar example of this is the way that, when one end of a stick is held in the hand and the other end tapped on the ground, the felt taps are, as it were, "projected" to the end of the stick—we feel the taps at the far end of the stick, not in the hand grasping the other end. An example of tactile projection familiar to automobile drivers is the way that bumps and slick spots in a road are felt in the tires of the automobile and not in the parts of the driver's body that are the sites of contact with the rest of the car.

Perhaps the most intriguing example of such sensory projection is given in experiments Paul Bach-y-Rita (1972) performed in developing prosthetic vision devices for the blind. The devices typically consist in a camera worn on the subject's head which sends low resolution video signals to a 16-by-16 or 20-by-20 array of tactile

stimulators worn on the subjects back. After only a few hours of training with the device, subjects could utilize the tactile stimulation fed to the surface of their skin by the camera to recognize distal faces and objects and read printed words that the camera was focused on.

Arguably, the trained subjects' experiences manifested many of the qualia of vision. Just as a person's tactile point of view can extend to the tip of a walking cane, the point of view of the trained subjects shifted from the sites of stimulation on their back to the point of view of the camera. Bach-y-Rita reports an occasion in which the zoom control of a subject's camera was activated without warning by an experimenter, causing the image in the camera to loom suddenly. The subject reacted by raising his arms to his head and lurching backwards (Bach-y-Rita, 1972, pp. 98-99). Although the tactile stimulus array was located low on the subject's back, the subject's behavior indicates that he located the percept high and in front of him. And in looming, the objects seen through the camera were seen by the subject as rapidly approaching him, indicating that the device was functioning as a distance sense for the trained subject.

Thus completes my brief catalog of the relevant and interesting spatial representation in sensory experience. I want to turn now to the topic of representationalism.

Representationalism about sensory experience gets its strongest foothold for any case in which we can point to a distinction between vehicle and content—between the experience itself and that which the experience is an experience of. The content/vehicle distinction gets a pretty good grip when we consider our sensory experiences of space. The experiences themselves, most will suppose, are (or at least token identical to) states of the brain. I can experience something as being far from me without the experience itself being far from me. I can experience something as being much larger than me without the experience itself being larger than me. The clearest way to explain the difference is by saying that experiences of spatial properties and relations represent those properties and relations and need not have those properties and relations in order to do so. The introspectible properties of spatial experiences are thus the contents of representations.

The psychosemantic question as applied to spatial representation is thus the question of how brain states come to represent spatial properties. How can a brain state that is not eight feet wide come to represent something as being eight feet wide? How can a brain state that is not a large triangular thing off to my right come to represent something as being a large triangular thing off to my right? I want now to describe procedural psychosemantics in general and then move on to describe procedural psychosemantics as applied to spatial representation.

An informational representation is about its causes and a procedural representation is about its effects. Prototypical instances of these two kinds of representation are neural representations in sensory cortex and motor cortex, respectively. Activations of motor representations in the brain eventuate in the sending of efferent signals to the effector organs. Sensory representations are activated by afferent signals from sensory organs. Informational representations are individuated by the flow of causation from the outside in, whereas procedural representations are individuated by the flow of causation from the inside out.

For a linguistic analogy to sensory and motor representations, consider the difference between observation statements and commands. The semantic content of an observation statement is its truth conditions and a reliable observer is one who is caused by the observed object to make true reports about the object. (Compare "The litmus paper is red" and the litmus paper's being red.) Instead of truth conditions, the semantic contents of commands are their success conditions. Successful commands serve to bring about the state of affairs that comprise their success conditions. (Compare "Clean up your room!" and a clean room.) Anscombe (1957, p. 56) makes a similar point illustrated by a the list read by a shopper in a grocery store, and the list made by a investigator following the shopper and noting what the shopper puts in the cart. The aptness of this analogy for present purposes is evident in the frequency with which efferent signals are equated with commands in the scientific and philosophical literature.

The proper psychosemantics for motor representations will be a procedural psychosemantics. The big question is, however, what could this have to do with spatial representation? I want now to consider how informational and procedural psychosemantics may each deal with the problem of spatial representation. I propose to pursue this by way of a thought experiment. The thought experiment is, without a doubt, both fanciful and highly artificial, as is typical of thought experiments. It is also, I think, quite illuminating, so I beg the readers' indulgence for the moment.

Imagine a creature living in an environment that consists of a smooth, flat, almost featureless, plane. The plane is topographically quite uninteresting, but it will allow our creature an extremely simple form of locomotion. The creature's means of locomotion consists in a pair of tank treads, thus I will hereafter call the creature "Tanky". I invite you to imagine the plane that Tanky traverses as a grid of squares. The width of each square is equal to the distance Tanky travels is a straight line if his tread turns exactly one full rotation. Just like a tank, Tanky can go straight forward and backward by turning both of his treads in concert and turn to the left or right by turning his treads in opposition. Most of the square regions in Tanky's planar world are devoid of anything noteworthy. However, a few are good places for Tanky to be (we can imagine the good squares being covered with a nutritious chemical) and a few are bad places for Tanky to be (we can imagine the bad squares being covered with a toxic chemical). Tanky is equipped with a sensory receptor on his underside so that when he is in a square, he can

detect whether it is nutritious, toxic, or neutral. A nutritious region may be eventually drained of its nutritious chemical by Tanky, and he then will leave to find a new one. The longer Tanky spends in a toxic region, the more deleterious it is for his health, thus it will be in his best interest to leave as soon as possible, and avoid future encounters with toxic regions.

We now have the gist of Tanky's environment and his basic survival goals. Everything that we need to know about Tanky's outer anatomy is exhausted by mentioning his two motor organs and his one chemoreceptor sensory organ. We must turn now to Tanky's inner anatomy. We are especially interested in finding out how it is that Tanky knows where he has been and figures out where he is going. The way Tanky knows whether he is in a region that is good, bad, or neutral is in virtue of afferent signals from his chemoreceptor. If we give Tanky the representational resources to have sensory experience and memory, then the psychosemantics will for his chemical representations will clearly be informational. Let us turn from his chemical representations to his spatial representations. Let us turn now to consider how Tanky might go about figuring out how far it is between, say, here and the next nutritious square (his next "square meal" one might say).

If Tanky has been traveling in a straight line, passed over a toxic square, and traveled through four neutral squares before encountering a nutritious square, how might

he come to know this? Let us consider two possible design solutions to this problem, one informational, and the other procedural.

The informational solution is as follows. First, in addition to Tanky's chemical receptor, he has sensory organs in his Tank treads that send afferent signals to Tanky's brain indicating whether the tread is turning forward or back, and how fast it is turning. Second, note that Tanky's motor commands to his effector organs are not accompanied by efference copies telling him what signal he sent. (An efference copy is a copy of the signal sent to the effector organs, but serves not as a trigger of the effectors but instead as a record of what commands were issued and when.)

On the procedural solution, however, Tanky's motor commands are accompanied by efference copies. Further, unlike the informational solution, on the procedural solution, Tanky would not have sensors in his tread indicating rate and direction of tread turning.

On both of the solutions we may assume that Tanky does not slip and slide in his travels, and no one ever picks him up and carries him to a different location. All of his travels across the plane are due to self-initiated turns of treads with good traction. On both solutions, then, Tanky has all of the efferent and afferent connections to his world he could need to give his memory and computational modules what they would need to compute his previous, current, and possible future locations. On both solutions, Tanky

has the representational resources for representing a nutritious and a toxic region as being, say, four tank tread distances away from each other.

The difference between the two solutions hinges on the different semantics of Tanky's spatial representations. On the informational solution, Tanky's spatial representations are triggered by afferent signals from his tread sensors. On the procedural solution, Tanky lacks tread sensors and his spatial representations are instead triggered by efference copies. On the informational solution, the spatial representations represent tank tread speed and direction in virtue of being caused by activity in the treads. On the procedural solution, the spatial representations represent tank tread speed and direction in virtue of causing activity in the treads.

The result of the Tanky thought experiment is to show that under certain conditions, procedural as well as informational representations can serve the needs of spatial representation. Informational representations do not have to be the only game in town. The thought experiment makes salient the possibility of the procedural solution to the problem of spatial representation. The thought experiment reveals the happenings in a possible world. The next question is: how is it done in the actual world? What solution did Mother Nature opt for us and our evolutionary relatives? Below I argue that the procedural solution is realized in actual cases. The points I raise below are conceptual and empirical considerations showing that motor differences—that is, efference differences—can give rise to differences in spatial experience.

The first consideration in favor of the behavioral constituency of perceptual space is a conceptual consideration due to George Pitcher (1971). Pitcher asks whether it is conceivable that the perceived directions from which sounds originate are not intrinsically tied to our dispositions to behave in ways appropriate to the perceived directions of the sound. As Pitcher points out, if it is so conceivable,

> . . .then it ought to be logically possible for someone to hear the direction from which a certain single sound (a bird-call, for example) is coming, and yet for him not to know in what direction he must point (or walk) if he is to point (or walk) in the direction from which the sound is coming, not to know in what direction he must look if he is to look in that direction, . . .and so on for all related abilities. (p. 189)

But clearly, such a possibility seems not to make sense. We just cannot imagine hearing a sound coming from a particular direction while simultaneously not knowing how to orient (or point, or walk. . .) toward the direction from which the sound seems to originate.

Similar considerations can be brought to bear on the visual apprehension of distributions of features in our visual fields. Imagine seeing a red stripe across your visual field such that its endpoints lie in the far left and right of your visual periphery. It makes little sense to suppose that one could perceive the endpoints of the stripe as occupying these regions of egocentric space without knowing how to orient to bring one or the other into one's fovea.

The case for the behavioral constituency of tactual-kinesthetic perception is even easier to make than for audition and vision. As Gareth Evans (1985a) points out, it is quite obvious that when a blind man uses his hand to perceive the spatial relations that parts of a chair bear to each other, the content of his experience is "partly determined by the disposition he has thereby exercised—for example, that if he moves his hand forward such-and-such a distance and to the right he will encounter the top part of a chair" (p. 389).

Our sensory awareness of spatial locations is not simply a matter of an egocentric coordinate system based on the arrangement of our own body parts, but instead a function of what we are disposed to do with those parts. As brought out nicely by Charles Taylor, up is not simply where our head is (as is evident when lying down), but instead a function of "how one would move and act in the [gravitational] field" (1979, p. 154).

Motoric output plays a constitutive role in the projective phenomena I had mentioned in connection with Bach-y-Rita's subjects. A necessary condition on being able to "see through" the camera-driven tactile array is that the subject be allowed to exert control over the inputs to the camera (Bach-y-Rita, 1972). Consider the following points in favor of this thesis.

Bach-y-Rita notes that the major portion of the first few hours of the subjects' training is occupied by learning the techniques of camera manipulation, including

controlling the operation of the zoom lens, the aperture and focus, and also the direction of the camera towards regions of the subjects' immediate environment (1972, pp. 3-4). Bach-y-Rita further notes that subjects with a high degree of manual dexterity acquire the ability to see through the prosthetic vision devices more quickly than those with a low degree of manual dexterity (1972, p. 7).

In discussing the case of the subject who raised his hands to protect his head in response to an unanticipated activation of the camera's zoom control, Bach-y-Rita writes:

> The startle response described above was obtained by a well-trained subject who was accustomed to have camera movement, zoom, and aperture under his control. (1972, p. 99)

Further, Bach-y-Rita writes, when subjects receive tactile array inputs from a static camera, they

> . . . report experiences in terms of feelings on the skin, but when they move the camera their reports are in terms of externally localized objects (1972, p. 99).

That the subjects be able to control the camera seems to be a necessary condition on the tactile stimuli on the back becoming transparent and giving rise to the devices functioning as a prosthetic distance modality. Indeed, this is the hypothesis advocated by Bach-y-Rita:

> [E]xternal localization of percepts depends critically on such movements and. . . a plausible hypothesis is that a translation of the input that is precisely correlated with self-generated movement is the necessary and sufficient condition for the experienced phenomena to be attributed to a stable outside world. Conversely, in the absence of such

> a correspondence, the origin is perceived as being within the observer (1972, pp. 99-101).

I must note, however, that I differ from Bach-y-Rita in that I deny that exercising motor control over the inputs is sufficient for the projection phenomena discussed. I wish only to claim that the motor control is necessary. Consider, in this regard, a datum presented by Bach-y-Rita: when a subject trained to see through the tactile array of the prosthetic vision device later scratches his back, he does not "see" the scratches (1972, p. 32). But, nonetheless, in scratching his back, he is in control of the tactile input. Thus, while necessary for instantiating the sensory projection essential to functioning as a distance sense, exerting motor control over one's sensory inputs seems not to be sufficient.

Plausibly, the necessity of motor control for sensory projection can be extended beyond the prosthetic vision case to the other instances of tactile projection I discussed above. Consider, for example, the case of the stick held in the hand. When one holds the stick in the hand and taps and rubs the end over textured surfaces, one feels the sensations at the end of the stick. Imagine, however, that the same information was delivered through the stick to the hand, but without your being able to control the rubbing and tapping of surfaces against the end of the stick. Imagine that you are blind-folded and strapped into a chair with your hands and arms restrained. A stick is placed in your hand and an experimenter taps and rubs the end of the stick with, say, a rough rock. It is not implausible to suppose that your tactile sensations would not project to the end of the stick, but instead would be perceived as occurring entirely within your hand.

If the considerations examined above are correct, then it follows that the natures of the qualities of phenomenal consciousness associated with stimuli in various modalities are determined not only by the nature of the information transduced by the nerve-endings in the sensory organs, but also by what types of subsequent motor activity that information is employed in. In the case of Bach-y-Rita's prosthetic vision experiments, the difference between perceiving skin stimulation as disturbances on the skin and seeing through the disturbances to a display of objects located in a three-dimensional egocentric space entailed differences in the kind and degree of motor control that the subjects were able to exert on the dynamics of the sensory input.

For another line of evidence linking efference and spatial qualia, consider what happens when one attempts to move one's eye when the eye muscles have been paralyzed. Gallistel (1980, p. 175) describes the following interesting phenomenon. A person may have their eye muscles paralyzed by curare. If a person with curare induced paralysis in her eyes attempts to shift her gaze to the left, the whole world will appear to her to have jumped to the left. This change in the visual scene is not due to any actual movement of the eye, nor has there been any change in the pattern of light on the retina. Gallistel (1980, p.175) writes that what has happened is that the typical expected association of efference copy and afference has been interrupted.

I interpret the case of the paralyzed eye as follows. Usually when one moves one's eyes to the left, the visual image moves to the right with respect to the eye. This typical association between the efferent output and afferent input is used to construct a representation of the world as a stable scene—even though the visual image is moving to the right with respect to the eye, the world is not seen as leaping to the right. But when the efferent command to shift is not accompanied by the afferent signal indicating the change of light on the retina (flowing to the right), the whole world is perceived not as stable, but instead as leaping in the direction indicated by the efferent command to shift to the left. Again, efference has had an effect on the introspectible properties of spatial experience.

We can easily imagine an analog to the paralyzed eye phenomenon in Tanky. In the procedural version of Tanky, he may construct a representation of his world based on expectations of which afferent signals would be paired with which efferent copies. We could paralyze Tanky's treads so that a motor command to go forward half a square is not accompanied by a change in the information presented to his chemoreceptor. A command to move forward, when not accompanied by the backward flow of the pattern of chemicals across the chemoreceptor surface would result in the perception of the world lurching forward.

A last empirical consideration I offer in favor of the proceduralist hypothesis is Held and Hein's (1963) famous kitten experiment. In this experiment a pair of kittens were placed in an environment controlled to make the visual input to each kitten as similar as possible, while allowing only one kitten to have control of the pattern of information coming into its eyes. The environment was an upright cylindrical volume with vertical stripes on the walls. In the center of the cylinder was a column. On top of the column was a horizontal bar that was attached at its center to a spindle. The horizontal bar could be turned around. Each kitten was attached to opposite ends of the bar. One kitten, the active kitten, was attached to the bar by means of a harness so that as the kitten walked around the cylindrical room, the bar would turn on the spindle. The other kitten, the passive kitten, was restrained in a gondola that hung from the other end of the bar. As the active kitten made the apparatus turn, the passive kitten was moved through the room at the same speed as the active kitten. When the kittens were not in the cylindrical room attached to the apparatus, they were both free to move about but only in the dark. After spending some time being reared in these conditions, the kittens were subjected to tests of their vision. The results were that only the active kitten developed normal visual perception. For example, only the active kitten avoided a visual cliff (that is, it avoided walking on a piece of glass suspended a considerable height over a surface below it)—a behavior typical of kittens its age.

We can easily imagine constructing Tanky style systems that would replicate results analogous to those in the kitten experiment. Imagine two procedural-style Tanky

creatures attached to the spindle, with one being allowed to actively move while the other was dragged along (with efferent signals to its treads blocked). Visual stripes on a wall would be replaced with chemical squares on the floor. Only the active tank would be able to develop certain efferent/afferent associations, associations like those discussed above in connection with the case of the paralyzed eye. The results of the kitten experiment show yet another connection between efference and perception and the Tanky analog to the kitten experiment helps show the point of entrance for a procedural psychosemantic interpretation of the kitten experiments. The passive tank would be subject to certain chemorecpetive inputs but be unable to correlate those with the effects that motor activity has on the pattern of inputs. Thus the passive tank would be unable to represent the spatial locations of the sources of the chemical stimuli.

One may question the relevance of the kitten experiment to the issue of the introspectible properties of sensory experience. It is not clear how we could have any idea as to what, if anything the kittens introspect. At least in the discussion of the paralyzed eye and the prosthetic vision devices, human subjects were involved, and we had their introspective reports as evidence of efferent-sensitive changes in spatial qualia. The kitten experiment is relevant, then, insofar as it may be considered as evidence that similar results would be obtained if humans instead of kittens were subjected to similar conditions. Such an experiment conducted on human children would obviously be unethical. So we can never be certain what would result from such an experiment. However, the kitten data, considered in light of the paralysis and prosthetic vision data,

lends further support to the proceduralist hypothesis about the nature of the representations in the experiences of spatial properties.

In this section I have offered a procedural psychosemantics as an alternative to the informational psychosemantics typically advocated by representationalists. I introduced the thought experiment of Tanky to urge the intuitive plausibility that a procedural psychosemantics could have anything to do with the way an organism represents spatial properties and relations. I discussed conceptual and empirical considerations that procedural representations actually are involved in some of our experiences of space. Note that I have not argued that all sensory representations must have procedural semantics. Nor do I think that all spatial representations necessarily have procedural semantics. I have argued that the informational story cannot be the whole story, and I have sketched a positive alternative—the procedural alternative—to handle those cases in which the informational story does not apply. I have established the plausibility that so-called egocentric representations of space are not mere motor commands, but in addition to being motor commands, they are also genuine representations of space.

2.3 Egocentric and Allocentric Non-spatial Representation

I turn now to show how the notion of perspectival representation applies to non-spatial representations. The point of showing this is two-fold. First, in the next chapter I

argue that the notion of perspective developed here will be sufficient to account for the subjectivity of conscious experience, especially as that notion arises in discussions of the knowledge argument against physicalism. Since that notion of subjectivity applies to experiences of non-spatial as well as spatial qualities, the notion of perspective developed here must be general enough to fit the bill. Second, the issues raised in this section are of significant interest apart from their contribution to discussions of consciousness. In particular, they constitute a continuation of the discussion of whether action-involving mental states may be representational. Of particular interest here is Akins' (1996) claim that the mental states involved in thermoreception are non-representational action-involving states. I argue that such states are indeed representational, but must be viewed as perspectival representations.

In a detailed examination of thermoperception (drawing on the work of Hensel (1982).), Akins (1996) finds grounds for questioning whether thermoreceptors allow the brain to represent features of the objective world (in this case, temperatures). Akins argues that in order for thermoperception to be in the business of representing temperature, three conditions must obtain. First, there must be constant correlation between receptor activity and temperature stimuli. Second, the activity of the receptors must preserve relevant structure of the stimuli (e.g. greater and lesser activity in the receptor must reflect greater and lesser temperature). Third, the sensory system must be servile in the sense that it does not embroider upon the information extracted from the environment. Akins argues that thermoperception fails all three criteria. While Akins'

critique provides an obstacle to viewing thermoreceptors as representing objective properties, I will argue that Akins' analysis leaves entirely open the possibility that these receptors serve to represent temperature in an egocentric way. Thus, on the account I favor, the contents of the deliverance of these sensory systems are less like reporting that the water is 5 degrees Celsius and more like reporting that the water is too cold for me. I agree with Akins that thermoperception is what Akins calls a "narcissistic" system. I depart from Akins in that I view thermoperception as representing narcissistic properties whereas Akins holds that thermoreception does not represent at all.

I focus on two aspects of Akins' discussion of thermoreceptors (cold receptors and warmth receptors) that lend to thinking of the human thermoperceptive system as producing perspectival representations of temperature. The first concerns the fact that thermoreceptors are not distributed across the skin in a uniform concentration, and that different concentrations of thermoreceptors give rise to different sensations given a particular skin temperature. The ratio of cold to warm receptors varies across different parts of the body. One result of this is that different parts of the body may have varying degrees of comfort for water at a given temperature. Water that feels comfortably tepid on the hands may feel shockingly cold when dumped over the top of the head. The second aspect of thermoreceptors that lend them to a perspectival view is that they have dynamic response functions. A thermoreceptor's response to a given temperature at a given time is in part a function of what its response at a previous instant was. This is evident in Berkeley's famous bucket example of the context sensitivity of temperature

perception. Prior to submerging your hands in a bucket tepid water, hold one hand in a bucket of ice and the other hand in hot water. The water in the bucket will feel hot to the previously chilled hand and much cooler to the previously heated hand. Both the differing concentrations of receptors and the dynamic response functions give rise to a many-to-one mapping of temperature sensations and temperatures. A sample of water of a given temperature will give rise to many different sensations depending on the concentrations of receptors and the level of their previous activity. These many-to-one mappings are arguably and plausibly part and parcel of the proper functioning of thermoreceptors. A given temperature may be more hazardous to tissue in one part of the body than another, and thus, a more sensitive alarm system may be accomplished by varying receptor concentrations. Dynamic responses may be adaptive since a rapid change of temperature can be damaging to tissue even if it occurs in a range of temperatures that would otherwise be harmless.

Akins sums up these aspects of thermoreceptors by describing them as "narcissistic": they are less concerned with how things are independently of the organism and more concerned with how things relate to the organism, thus echoing the human narcissist's favorite question: "so what does this have to do with me?". Akins takes the narcissism of thermoperception to count against the claim that thermorecptive sensory systems represent at all. I favor the alternative interpretation that thermoperception represents temperature, albeit in a perspectival way. A given temperature sensation does

not just represent a temperature of a region on or near the skin but represents temperatures as being of varying degrees of hazard or harmlessness to the subject's tissues. Thermoreceptors don't simply represent temperature, but include in the representational contents of their outputs relations that the temperatures bear to the representing subject, much in the way that retinocentric representations of spatial locations represent locations defined relative to the subject.

Arguably, perspectival representation may be found in examples beyond spatial and thermal perception. Our detection of chemicals in olfactory and gustatory senses may not be in the job of simply representing the presence of a certain chemical but also representing the chemical as noxious or poisonous or nutritious. But these are properties that can only be defined in relation to the organism: one man's meat is another man's poison and all that. Thus any system that has the function of causally covarying with such properties thereby produces perspectival representations of chemical concentrations.

Akins offers arguments against the kind of move I favor, but I find her arguments wanting. The proposal under consideration may be described as Akins describes it: the proposal that thermoreceptors do not represent temperature but instead temperature-involving narcissistic properties. Thus, the output of a thermoreceptor in response to a given temperature does not represent a given temperature per se but instead whether the given temperature is, e.g., too hot, too cold, or just right. The property of being too hot cannot be defined independently of answering the question "too hot for whom?" and the

subject relativity of such a property is what makes it narcissistic. Might thermoreceptors have the function of detecting narcissistic properties? Akins' objects to the proposal that thermoreception represents narcissistic properties on the grounds that the proposal depends on accepting the Detection Thesis: the claim that "each and every sensory system functions to detect properties" (p. 360). Akins then gives two reasons for doubting the Detection Thesis: the first is that the thesis is overly strong given the relatively small amount of evidence regarding sensory function collected to date. The second is that in at least one case—the case of proprioception—interpreting sensory activity as having the function of detecting is unhelpful and unenlightening.

I propose to grant Akins the falsity of the Detection Thesis, since such a concession leaves unscathed the proposal I favor. For convenience I will call the proposal "the narcissistic representation proposal"—the proposal that thermoperception represents narcissistic properties. Akins misconstrues the logic of the situation in asserting that the narcissistic representation proposal depends on the Detection Thesis. Contra Akins, the Detection Thesis is not a necessary condition on the truth of the narcissistic representation proposal. The narcissistic representation proposal plausibly has as a necessary condition the truth of the thesis that at least one sensory system functions to detect properties. But it is not at all obvious how it could have as a necessary condition the claim that "each and every" sensory system functions to detect properties. And for my immediate purpose, the purpose of establishing the plausibility of non-spatial

representations that nonetheless have pictorial perspective, the detection thesis may be disregarded as irrelevant.

Chapter 3: The Subjectivity of Conscious Experience

3.0 Introduction

Conscious experiences are supposed by many to be subjective in the sense of being perspectival or from a point-of-view (see, for example Nagel 1974, 1986 and Tye 1995). Allegedly, the subjectivity of consciousness is beyond the grasp of science, which is objective (Nagel 1986). This claim is supposed to follow from the famous Nagel-Jackson knowledge argument. If you wanted a scientific understanding of consciousness, how would you solve the problem of subjectivity? One strategy would be to deny that consciousness is really subjective. Another might be to deny that science is really objective. Yet a third, and the one I favor, conserves both the subjectivity of experience and the objectivity of science. Drawing on notions developed in the previous chapters, I present a way of looking at subjectivity and objectivity that allows for an objective, scientific—indeed, neuroscientific—understanding of the subjectivity of conscious experience. I critique recent representational accounts of experience—such as those advocated by Georges Rey (1997), Michael Tye (1995), and William Lycan (1996)—that explicate subjectivity in terms of indexicals. I raise problems with the indexical account that hinge on its treatment of the knowledge argument. I then sketch a positive account of subjectivity, one that is both physicalistic and representational, but does not fall prey to

the problems that beset the indexical account. This account explicates the subjectivity of experience in terms of the notion of pictorial perspective developed in the previous chapter.

3.1 The Paradox of Subjectivity of Conscious Experience

In chapter 1 I introduced the paradoxes of subjectivity for both secondary qualities and conscious experience. In that chapter I presented a solution to the paradox for secondary qualities. Here I address the paradox as it concerns consciousness. Recall that the paradox is as follows.

> Premise 1. What it's like to be a bat is subjective
> Premise 2. What it's like to be a bat is just the bat's being in a particular brain state
> Premise 3. Whether the bat is in a particular brain state is entirely objective
> Conclusion: What it's like to be a bat is objective

Michael Tye (1995, pp. 56-62) writes of a "paradox of phenomenal consciousness" that can be shown to be similar to the paradox I am concerned with. Tye writes:

> (7) Phenomenal states are perspectivally subjective.
> (8) If phenomenal states are perspectivally subjective, then they are neither identical with, nor realized by objective physical types.
> (9) If phenomenal states are neither identical with, nor realized by, objective physical types, then they are not even broadly physical states.
> (10) If phenomenal states are not even broadly physical states, then they are causally irrelevant.
> Therefore,
> (11) phenomenal states are causally irrelevant.

> Unfortunately, (11) is clearly false. What makes for a paradox is that all or the premises seem clearly true, once we reflect on them in the context of the problems presented in this chapter. And (11) follows from (7) through (10) via the rules of formal logic. So, we should be deeply perplexed. (p. 62)

This can be restated to more closely resemble the paradox of subjectivity as I have formulated it. First, note that the contraposition of (10) is:

> (10') if phenomenal states are causally relevant, then phenomenal states are broadly physical states.

Next note that the contraposition of (8) is:

> (8'): If phenomenal states are either identical with, or realized by objective physical types, then they are not perspectivally subjective

Let Tye's assertion of the falsity of (11) be:

> (12) phenomenal states are causally relevant.

We get from Tye's paradox, to mine as follows.

> (7) Phenomenal states are perspectivally subjective.
> (13) phenomenal states are causally relevant.
> (10') if phenomenal states are causally relevant, then phenomenal states are broadly physical states (that is, they are either identical with, or realized by objective physical types)
> (8') If phenomenal states are either identical with, or realized by objective physical types, then they are not perspectivally subjective
> (13) Phenomenal states are not perspectivally subjective (from 12, 10', and 8')

Again, an attempt to physicalize subjectivity (this time by way of causal relevance) leads to a conclusion that states of phenomenal consciousness both are and are not subjective. The conflict between thinking of experience as subjective and its physical basis as

objective reaches an especially potent head in the knowledge argument against physicalism, to which I know turn.

3.2 Indexical Knowledge

The gist of the knowledge argument is as follows. Mary has never seen red. She nonetheless knows all the physical facts. Upon seeing red for the first time, Mary learns something new: what it is like to see red. Prior to seeing red, Mary knew all the physical facts; thus in learning a new fact upon having a red experience, Mary learns a non-physical fact. Thus the subjectivity of experience is non-physical (Jackson 1982).

The indexical account of subjectivity is due to physicalists such as Lycan (1996) and Tye (1995). To take one instance as representative, I sketch Lycan's account. Experiences are representations. When I have a conscious experience of a fire engine being bright red and six feet away from me, the experience itself is neither bright red nor six feet away from me. The experience itself is a state of my nervous system that represents the fire engine as being bright red and six feet away from me. When I introspect my experience, I form a second-order representation of the first-order representation of the fire engine. Other people may form syntactically similar second-order representations, but those representations will be about their first-order states, not mine. The crucial analogy here is to the use of indexicals in speech. When I say "my leg

hurts" I am referring to my leg, and only I can refer to my leg by using that utterance. You may use a syntactically similar construction: you may utter the words "my leg hurts", but in doing so, you would be representing your leg, not mine. Analogously, only I can represent my first-order states by the introspective application of self-referential indexical concepts. And this, according to Lycan, is the ultimate explication of subjectivity.

The indexical response to the knowledge argument is one of the many responses to the knowledge argument that hinge on the notion that there may be multiple modes of presentation of a single physical fact. This kind of response grants that Mary learns something new but only in the sense of learning to apply a new mode of presentation to an old fact. This kind of defense of physicalism falls prey to the objection that in learning to apply a new mode of presentation to an old fact, the subject learns a new fact, namely, that the new mode of presentation applies to the old fact. Given the presupposition that the subject already knew all of the physical facts, this new fact must be non-physical (Alter 1998).

Another problem with the indexical response is that it mistakenly makes numerical differences sufficient for subjective differences. To see why this is a bad thing, consider the following. Suppose that while Mary does not know what it is like to see red, Cheri, Mary's color-sighted colleague does know what it is like to see red. Upon seeing red for the first time, not only does Mary learn what it is like to see red, she learns

what it is like to be Cheri. If Mary and Cheri were physical and experiential doppelgangers (though numerically distinct individuals) they could each know what it is like to be the other person, regardless of whether their numerical non-identity entails divergence of the contents of their indexical thoughts. As such, then, Mary and Cheri would be subjectively identical, in spite of being indexically distinct. Thus indexicality is inadequate to account for subjectivity.

This point about what Mary and Cheri know needs to be stated carefully, for the point is consistent with the fact that for all Mary knows, there is nothing at all it is like to be Cheri. The point is that if there is something it is like to be Cheri when Cheri is having experience of kind *X*, and if Mary knows what it is like to have experience of kind *X*, then Mary knows what it is like to be Cheri. This is analogous to knowing what the Mona Lisa looks like without knowing that one has ever seen the Mona Lisa or that there even is such a painting as the Mona Lisa.

Another way to state the point that I am after is that if Mary and Cheri where physical and experiential doppelgangers (though numerically distinct individuals) they could each know what it is like to be the other person, regardless of whether their numerical non-identity entails divergence of the contents of their indexical thoughts. What I'm denying is that any two individuals necessarily differ subjectively. This does not require a blanket reliance on a problematic distinction between intrinsic and extrinsic qualities. All it requires is that things can be identical in some respects while differing in

others. Consider the question of whether any two books necessarily differ in the plots of their stories. Arguably, they do not: two copies of *Moby Dick* may be identical with respect to their plots. Analogously, intuition tells us that two individuals that are spatio-temporally distinct may nonetheless be identical with respect to the subjective aspects of their experiences.

3.3 Pictorial perspective and the knowledge argument

In the chapter 2 I have proposed that some mental representations exhibit pictorial perspective. In the present chapter I need to tie this into consciousness. Do states of consciousness possess this kind of perspective? And what about the so called "knowledge argument" that has figured heavily in discussions of the subjectivity of consciousness?

Regarding whether conscious states exhibit this kind of perspective, the answer seems a resounding "yes". The thermoperception examples are all examples of conscious sensations that vary independently of actual temperature: what enters into sensation includes relations of the temperature to states of the subject. Water of a given temperature may feel colder on the head than on the hands. Likewise, the remarks about the phenomenology of visual experience lead naturally to finding pictorial perspective in

conscious representation. My percept of seeing a house differs depending on where I am standing. It depends on my literal point of view.

On the alternate response to the knowledge argument that I favor, the first premise of the argument is false: it is false that the subject could know all the physical facts without having an experience of red. I favor the view that there are both objective and subjective physical facts. What a subject can learn only by having an experience of red is a subjective, yet nonetheless wholly physical fact. Thus, while Mary may know all the *objective* physical facts without seeing red, what is left out until she learns what it is like to see red is a set of *subjective* yet nonetheless physical facts. Similar views have been defended elsewhere (see for example Deutsch (unpublished) and McGinn (1991). Indeed, Nagel's (1974) original aim in drawing attention to what it is like to be a bat was not to defeat physicalism, but instead to argue that the objective did not exhaust the physical. However, previous accounts along these lines give very little detail about what such subjective physical facts might consist in or why such facts deserve to be regarded as physical. By accounting for subjectivity in terms of perspectival mental representation, I provide remedies to such problems.

The account of pictorial perspective described above renders intelligible the compatibility of physicalism and subjectivity. It allows us to see how a physicalistic framework can tolerate first, physical properties that depend for their existence on representations and second, physical properties that can be represented only by the

representations that they depend on. Thus, such physical facts are subjective in the classical senses of being first, mind-dependent, and second, knowable only by a restricted mode of access. These notions may be briefly characterized by reference to imagistic representations.

The subject-dependence involved may be sketched as follows. What the image represents depends for its existence on the process of its representation. Precisely what is represented cannot be characterized independently of specifying the point of view of the representing subject. For example, part of what is represented is what the object looks like from one location as opposed to another. The notion of subject-dependence involved here may be conveyed by the following example. Consider a pictorial representation of a complex object like Mount Everest. Imagine that portions of the surface of the mountain are painted black so that from a particular point of view only black regions of the mountain could be seen, but from any other point of view, many of the mountain's non-black surfaces can be seen. Consider the set of black regions of the mountain. What is it that unifies those regions as a set? What is common to all and only those regions? The point of view occupied by a viewer—a generator of pictorial representations—is the unifying essence of those particular regions. It is in this sense, then, that the things that are represented depend on being represented. Of course, there is a sense in which those regions would exist even if no one were to represent them. However, in specifying the set comprised of all and only the spatial regions captured in the image, one does not carve nature at its joints, but instead carves nature into a gerrymandered collection of items that

would be of no interest apart from their involvement in a particular representation. That much of neural representation is concerned with such gerrymandered properties should not come as an enormous surprise. For instance, it makes sense that an animal's chemoreceptors would be less interested in carving nature into the periodic table of elements and more interested in carving nature into the nutrients and the poisons—categories that make no sense apart from the needs of that organism.

The restricted epistemic access involved may be sketched in terms of imagistic representations as follows. What is represented by an image can only be represented, without addition or deletion, by an image. Even a string a numerals coding a bitmap for an image does not have all and only the representational content of the image. The numeral string constitutes, in part, a recipe for constructing an image, and in doing so, it has content that the image itself lacks. The old saw about a picture's being worth a thousand words is false: a picture is worth no number of words. This point cannot be adequately argued for here, but suffice it to suggest that it is not obviously incompatible with physicalism that there are properties represented in sensory experience that may only be represented in sensory experience. Of course, *part* of what is represented in olfaction may be conveyed in some other mode of representation like the phrase: "your perfume smells like vanilla and roses". But the suggestion that *all and only* what is represented in olfactory experience can only be represented in olfactory experience is an entirely physical possibility, if not a physical actuality. That experience is perspectival in this sense allows us to conceive of physical facts that may be knowable only by a

restricted mode of access, that is, physical facts that may only be represented by specific sensory experiences.

This latter point may be further unpacked in terms of the theory of mental representation mentioned earlier. On that theory, the process of representation is constituted by certain causal relations obtaining between a state of an organism (the representation itself) and an environmental state (that which is represented). Objective facts may be cashed out in terms of environmental states that are capable of entering into many kinds of causal interactions in addition to those in virtue of which they are represented. So, for instance, not only does water causally interact with my representation of water, water also causally interacts with the mountains it erodes, the crops that it irrigates, and so on. Many causal chains lead to water thus providing for multiple routes of epistemic access to water—multiple ways of knowing about water. In contrast, subjective facts may be cashed out in terms of things that are able to enter into causal interactions only with the mental states that represent them. Such states do not admit of multiple chains of causal interactions leading to them and thus do not permit of multiple ways of being known. In causally interacting only with representations of them, such states are knowable only in virtue of being represented, as, for example, under the description "that which I am experiencing now" and cannot be known through other means. As such they are subjective facts, knowable only through restricted modes of access. But we may be confident that such facts are nonetheless wholly physical, since their possibility is provided for by a wholly physicalistic account of mental

representation. (I further address these epistemological features of subjectivity in chapter 5, in a discussion of the connection between subjectivity and private access.)

Thus, armed with an account of pictorially perspectival mental representations explicated in terms of causal relations, the physicalist may fend off concerns stemming from the knowledge argument. These remarks are far too brief to *establish* that thinking of subjectivity in terms of pictorial perspective renders the knowledge argument against physicalism ineffectual. However, these remarks do show, at a minimum, that it is not obvious that abandoning the indexical account of subjectivity leaves the physicalist defenseless in the face of the knowledge argument.

Chapter 4: Objectivity and Space

4.0 Introduction

In previous chapters, the primary topic has been the distinction between the objective and the subjective. Perhaps surprisingly, the topic of spatial representation has surfaced on several occasions. Surprising or not, the philosophical connection between objectivity and space has a long lineage.

For several centuries and in many areas of philosophy various species of the distinction between the objective and the subjective are expressed in a spatial idiom. Philosophers alternately worry about and shrug off the problem of the external world. They wonder whether anything exists outside of the mind. Metaphorical articulations of the notions of subjectivity and objectivity exploit the spatial idiom of seeing things from a point of view. What can be seen from only one point of view is more subjective and less objective that what can be seen from any point of view. The maximally objective view is, to use Thomas Nagel's (1986) phrase, the view from nowhere.

Space and objectivity are associated in doctrine as well idiom. For Kant, space is the form of outer sense. The association of space and objectivity is particularly strong in

Cartesian dualism whereby whatever is objective is physical, that it, has spatial magnitude, whereby whatever is subjective is nonphysical, that is, lacking in spatial magnitude.

Philosophers like Thomas Nagel (1986) have worried that the subjectivity of conscious experience bars the possibility of giving a physicalistic explanation of consciousness. Colin McGinn (1995) explicitly targets the non-spatial aspects of consciousness as rendering them recalcitrant to physicalization. The classical distinction between primary and secondary properties, i.e., the distinction between objective properties of objects and those "in the eye of the beholder" properties fell along spatial/non-spatial lines. Primary properties, that is, objective properties, of objects were cashed out in terms of the occupation of and movement through space (Evans 1985, pp. 268-281). In the philosophy of mathematics, there is a pervasive unease about attempts to cash out the objectivity of mathematical knowledge in terms of reference to non-physical objects (Benacerraf (1965) and (1973), Dummett (1975), Field (1989), Quine (1980), and Wrenn (1998).)

The philosophical association of space and objectivity takes on an especially acute focus in the work of Peter Strawson (1959). Strawson argues that our concepts of material bodies, that is, space-occupying particulars, occupy a basic role in our conceptual scheme, meaning that it is a condition on admitting any kind of particular into our ontology that we know how such particulars relate to material bodies (Strawson 1959,

chapter 1). Strawson is especially interested in arguing that this condition holds of objective particulars, that is, particulars that are conceived to be neither states of ourselves nor dependent on states of ourselves (Strawson 1959, chapter 2). (These are particulars conceived of in terms of what I have called "metaphysical objectivity" in my chapter 1.) The general form of Strawson's arguments to this conclusion proceeds by attempting to show that any creature incapable of representing spatial properties would thereby be incapable of representing objective particulars. Strawson's student Evans (1985) presents different arguments for similar conclusions. One way of reading Strawson and Evans' arguments is as arguing for the conclusion that any creature able to have epistemically objective representations must be capable of representing objects as having spatial properties. On another reading, the conclusion of their arguments are even stronger, namely, that any epistemically objective representations must represent objects as having spatial properties.[3] If this stronger conclusion is true, then the account of

---

[3] It must be noted that the textual evidence for the stronger reading is not univocal. For example, Strawson asks "is the status of material bodies as basic particulars a necessary condition of knowledge of objective particulars?" (1959, p. 61) indicating that his aim is to show only that material bodies are basic elements in acquiring knowledge of objective particulars, not that they are the only particulars about which such knowledge may be had. Even though Strawson here may be read as contradicting what I am calling the strong reading, it is not clear that his arguments are in strict accord with Strawson's

objectivity presented in previous chapters is in serious jeopardy, since in chapter 2, I argued that objective/subjective distinction cross-classifies the spatial/non-spatial distinction. Strawson and Evans' arguments have relevance to the current project aside from the possibility that this stronger thesis is true (or that they even argued for it). Even the truth of the weaker thesis poses a threat to my project. In chapters 2 and 3 I argued for an account of objective representations such that relatively simple creatures may have them, including creatures too simple to represent spatial properties. Because of these possible threats, in this chapter I consider and counter Strawson and Evans' arguments that objectivity requires space.

I begin with some brief remarks about space.

Real space and quasi-space.

I need first to distinguish real space from mere quasi-spaces. I cannot provide anything like definitions by which to guide the distinction, but perhaps the following will suffice. An object that changes location from London to Paris has moved through real space. An object that changes from red to orange to yellow may have moved through a quasi-space—a "color space"—but need not have moved through real space to do so. This is because dimensionality, though necessary, is not sufficient condition for being real space.

---

statement. And even if they were, it is another question entirely whether Evans arguments are so restricted. I will have more to say about these points below.

Colors and sounds may vary along dimensions (e.g., hue and pitch) but these are not genuine spatial dimensions as those involved in varying your location from Chicago to New York. The distinction between dimensions that do and dimensions that do not satisfy the sufficient conditions for being space is the distinction between real space and merely quasi-spaces. I hereafter will simply call real space "space".

Spatial Properties and Spatial Relations.

Describing objects as square or spherical is to predicate spatial properties of them. Describing an object as to the right of or three feet away from another is to predicate real spatial relations of those objects. I do not have definitions of spatial properties and relations and must make do with examples such as those given so far. Describing an object as the farthest pyramid from the Taj Mahal is to predicate spatial properties and relations of that object. Describing an object as a loud stinky blue thing is not to predicate spatial properties of that object.

Not all experiences are spatial.

Describing one sound as louder than or a higher pitch than another is insufficient to attribute real spatial properties or relations to the things heard. One may supplement one's theory of the world with theorems whereby one may infer that louder sounds are closer than quieter ones, or that higher pitched sounds are moving toward you and lower pitched sounds moving away. But, the incorporation of such theorems are optional, I will suppose. This much seems beyond doubt: we may experience things without

experiencing them as having spatial properties or relations. Olfaction delivers many of what I regard to be non-spatial experiences. One may smell one odor to be sweet and another to be pungent without smelling either to be anywhere at all. In much of this paper, I will follow Strawson and Evans in supposing that auditory experience constitutes a source of non-spatial experiences.

Thus completes my brief preliminary discussion of space. It is the sense of "space" just indicated, then, that I will consider the thesis that space is a requirement of objectivity.

What sense then, can we make of the claim that objectivity requires space? I offer three interpretations of the thesis at stake. Space will count as a requirement of objectivity only if at least one of the following sentences are added as theorems to the predicational theory of objectivity:

> SO1. A sentence or attitude is epistemically objective only if it contains a predicate that names a spatial property or relation.
>
> SO2. A property or relation is metaphysically objective only if it is a spatial property or relation
>
> SO3. An object is metaphysically objective only if it has spatial properties or bears spatial relations to some other object(s).

Below I argue that Evans' and Strawson's arguments fail to establish the need to incorporate any of the three propositions, SO1-SO3, into my theory of objectivity.

There is no doubt that this short list fails to exhaust the multifarious ways one could disambiguate the claim that objectivity requires space. I justify the meager length of the list on the following grounds. First, adding more items on the list while giving them the attention that they deserve arguments that follow would bring this discussion to an excessive length. Second, I justify the inclusion of SO1 on the grounds that it is the most relevant to concerns arising over the notion of allocentric cognitive representations discussed in previous chapters. If something like SO1 is true, then the notion of allocentric representation cannot be extended to apply to representations that do not represent objects as shaped or located in space. Third, I justify the inclusion of SO2 and SO3 on the grounds that these are very close, if not identical, to theses that Strawson and Evans argue for under the heading of a defense of the spatial requirements of objectivity. More specifically, one of Evans' arguments concerns SO2, while two of Strawson's and another of Evans' concern SO3.

Note that given the predicational theory of objectivity, SO2 entails and is entailed by SO1. SO3, however, is logically independent of SO1. Strawson's and Evans' arguments for SO2 and SO3, if sound, would require the inclusion of SO1, SO2, and SO3 as theorems of my predicational theory of objectivity. If my arguments below are sound, however, then Strawson and Evans do not supply compelling reasons for the inclusion of SO1, SO2, and SO3 in the predicational theory.

Below I examine and critique two arguments of Strawson's that objects (or particulars, to use the term that Strawson favors) must have spatial properties and relations in order to be metaphysically objective. I shall call these arguments "The Reidentifiability Argument" and "The Elsewhere Argument". Evans offers an argument regarding the metaphysical objectivity of objects and I follow Strawson (1980) in calling the argument "The Simultaneity Argument". Evans also offers an argument that may be construed as an argument for the spatiality of any metaphysically objective properties. I follow Strawson in calling this latter argument "The Causal Ground Argument".

4.1 Strawson's Reidentifiability Argument

At the heart of all four arguments is Strawson's thought experiment from his "Sounds" chapter (chapter 2) of *Individuals*. In this thought experiment Strawson invites his readers to attempt to imagine a subject that does not experience things as having spatial properties or relations. The point of this thought experiment is to see if it is conceivable that such a subject be able to grasp the concept of objectivity. According to Strawson,

> the question we are to consider, then, is this: Could a being whose experience was purely auditory have a conceptual scheme which provided for objective particulars? (p. 66)

(For ease of exposition, Evans called Strawson's imagined subject "Hero". I follow this practice and also shall, for ease of exposition, call the imagined purely auditory world "Auditoria".)

Strawson sets out to see if Hero can "make sense of" and "have a use for" a concept of objective particulars (1959, p. 69). Toward this end, Strawson sets out to see if Hero can make sense of the notion of particulars, postponing their objectivity for the moment. According to Strawson, in order to get a decent concept of particulars in Auditoria, one must (i) get identifiable, in the sense of distinguishable, sound-particulars (1959, pp. 69-70) and (ii) get identifiable, in the sense of reidentifiable, sound-particulars (1959, pp. 70). For Strawson, having an auditory perception is sufficient for the identification of a sound particular, and the auditory experience of continuity and discontinuity is sufficient for distinguishing sound particulars. A *C#* ("*C#*" names a universal here) that plays (gets instantiated for a duration), stops, then plays again, gives an example of two distinguishable tokens of the same type, and if the note played did not stop and start again, there would be just be one token. But for the reidentification of sound particulars, more than continuity and discontinuity of sounds is needed. According to Strawson, spatial criteria are needed. Below I describe why Strawson thinks space is needed for reidentification, but for now I describe what Strawson thinks reidentification is. According to Strawson, a particular is reidentified if and only if it is perceived for some continuous period that ceases, and perceived some second time and identified by

the perceiver as the numerically same particular perceived earlier. Thus a particular is reidentifiable only if it can be perceived twice.

Reidentification lies on the path from objectivity to spatiality in Strawson's argument. Objective particulars must be able to exist unperceived. Strawson argues further that any particular that exists unperceived must be reidentifiable, which means, if it is perceived for some continuous period that ceases, it must be able to be perceived some second time and identified by the perceiver as the same particular perceived earlier. Strawson also argues that the only criteria by which one may sensibly regard some perceived thing as the same as some particular earlier perceived are spatial criteria.

There are, then, two key stages to Strawson's argument, both of which I call into question. The first is the argument from the metaphysical objectivity of particulars to their reidentifiability. The second stage is the argument from reidentifiability of particulars to the necessary employment of spatial criteria for their reidentification.

Strawson argues that objectivity entails particular reidentifiability because objectivity entails that there are things that exist independently of whether one perceives them. And this in turn entails the logical possibility of perceiving something at two different disjoint times that persists unperceived between those two times. Thus, if it is possible for x to exist while you are perceiving it and it is possible for x to exist while you are not perceiving it, then it is possible for you to see x at time t, not see x at t+1, and see

x—the numerically same *x*—again at *t*+2. And, if you can see *x* twice, then you could identify *x* twice, that is, identify *x* and then reidentify *x*.

Is Strawson right that objectivity entails particular reidentifiability? Is every objective particular that you can perceive an object that you can perceive twice? No and no. I offer that it is simply false that every metaphysically objective particular that you can perceive at some time *t* must be perceivable at some time *t*+*n*.

Consider a particular time slice of an event or process (a slice thick enough to be perceived at least once). This seems like a particular that you can perceive only once, but nonetheless might have existed even if no one perceived it (or represented it in any other way). When I booted up my computer this morning, I perceived the beginning of that process—I perceived the particular initial time slice of my computer's booting up this morning. And, I suppose, that particular time slice may have existed unperceived. But that particular time slice is not something I can perceive a second time. The moment has passed, alas. To multiply examples, consider also extremely short lived particulars, like particular explosions or particular flashes of lightning. Again, they seem plausibly metaphysically objective without being reidentifiable.

Chase Wrenn (personal communication) offers as a possibly perceived but not reidentifiable class of particulars the fictional (we assume) class of particles called "nihilons". Nihilons exist until they are finished being perceived for the first (and only)

time. A nihilon may exist unperceived, but once perceived, its existence is terminated at the end of an episode of perceiving it. Any reidentified entity would not be a nihilon. Since Strawson's inquiry concerns the structure of our conceptual scheme, and since nihilons are clearly conceivable, then the conceivability of objective particulars seems not to require their reidentifiability. Strawson is committed to the claim that for any object that we can conceive of perceiving and also conceive of existing unperceived, we can further conceive of perceiving on more than one occasion. Thus, he is committed to the claim that nihilons are inconceivable. But clearly they are conceivable. At least Chase Wrenn and I can conceive of them.

Counterexamples to the claim that objectivity entails particular reidentifiability include (i) particular time slices of long duration events and processes, (ii) short duration events and processes, and (iii) nihilons.

Strawson takes himself to have shown that objectivity entails reidentifiability. His next move is to show that trying to get reidentifiability into Auditoria will require treating one of the dimensions of Auditoria as an analog to spatial dimensions in our conceptual scheme. According to Strawson, some non-temporal dimension in Auditoria must be sought to go proxy for the absent spatial dimensions. This proxy dimension will provide subsidized housing for unperceived yet enduring sound-particulars.

Auditory candidates for this dimension include timbre, pitch, and volume. Strawson dismisses timbre immediately for its lack of obvious systematic ordering of different timbres. (I suspect this to be due to the fact that any plausible ordering scheme for timbres will be multidimensional.) Strawson prefers pitch. So be it. Experiences in Auditoria will have the following structure. Items in Auditoria are sounds and sound sequences like pieces of music. Items are accompanied by continuous back-ground noise known as the Master Sound. The pitch of the Master Sound is going to be Auditoria's pseudo-spatial dimension. The "location" of a particular sound or sound sequence is that pitch of the Master Sound that is contemporaneous with the sound sequence instance. Suppose that particular sound sequence is a particular playing of *Ode to Joy* (a dated occurrence or tokening of the song type/universal *Ode to Joy*). This instance of *Ode to Joy* is heard by Hero over a finite duration. At a particular instance, the pitch of the master sound is, say, *C#*. Imagine hearing the pitch of the Master Sound increase while *Ode to Joy*'s volume decreases, and eventually, at some pitch of the Master Sound, *Ode to Joy's* volume is inaudibly low. Imagine further that the increasing pitch of the Master Sound, accompanied by the decreasing volumes of *Ode to Joy* is also accompanied by the increasing volume of some other sound sequence instance—*Jesu: Joy of Man's Desiring*. All this is reversible too: as the pitch decreases, returning to *C#, Jesu: Joy of Man's Desiring* gradually quiets down while *Ode to Joy's* thunder swells. Thus, during duration *d* in which the pitch of the Master Sound is going up and down, *Ode to Joy* is maximally audible while contemporaneous with one pitch of the master sound, whereas *Jesu: Joy of Man's Desiring* is maximally audible while contemporaneous with a different pitch of the

master-sound. The intuition being urged here is that these two different sound sequence instances, *Ode* and *Jesu*, exist during the same duration *d*, but at different "locations" i.e., different pitches of the Master Sound.

At least as far as Strawson is concerned, we have now accumulated the minimal ingredients for citizens of Auditoria to have a distinction between numerical and qualitative identity. Qualitative identity/distinctness is easy, and needs no further comment. Numerically distinct tokens of the same type might be conceived of as follows. If, at some pitch of the master sound *L*, *Jesu: Joy of Man's Desiring* is maximally audible, and the pitch rises to some other pitch *L'*, while *Jesu* gets quieter, and then promptly changed back to *L* while *Jesu* returns to its maximal volume, then we may suppose that the same token has been heard at different times. If, however, we go from *L* past *L'* to *L"* while *Jesu* gets quieter and then louder, we may suppose that we have come across a different (numerically distinct) tokens of the *Jesu* sound sequence type.

One might object here that in this imagined case it is hard to pretend that there is anything like a fact of the matter about such questions of "numerical identity". I interpret this objection as the worry that what constitutes the conditions of numerical identity for sound particulars is more a matter of a judgment-call on our part and less a matter of a discoverable and mind-independent fact. I suspect that Strawson might respond by agreeing that it is indeed more a judgment-call on our part. But, I think Strawson would add, the point of the project is to investigate what kinds of experiences would allow for

us to have uses for and make sense of certain kinds of judgment-calls. (This is the sense I make of Strawson's use of the phrases "make sense of" and "have a use for" with regards to certain concepts that may or may not be admitted in Hero's repertoire (1959, p. 69).) Suppose that we introduce by stipulation the term "blorg". The proper use of the term, since stipulated, will constitute a judgment-call on our part. But only if we have certain kinds of experiences will we be able to have a use for certain stipulations about the term's applicability to experience.

Strawson's question regarding the numerical identity criteria of sound particulars is the following. Could any criteria of numerical identity—and thus reidentifiability—be stipulated that would allow for the application of the concept of objective particulars to experiences that did not consist of representations of objects as having spatial properties and relations? Strawson answers "No" but I answer "Yes". Criteria of reidentifiability need not be spatial. This can be shown by showing that criteria of reidentifiability can be spelled out with out appeal to one of the necessary conditions on spatiality, namely, dimensionality.

The Master Sound need not be a set of sounds ordered along some dimension like pitch or volume. The Master Sound could be a set of the following sounds: the sound of a washing machine, the sound of a saxophone, and the sound of a baby crying. If at one time, Hero hears *Jesu* accompanied by the baby cry and at another time, *Jesu* accompanied by the sound of washing machine, those would count as two different

instances of the same sound type. If instead, both times Hero heard *Jesu* being played along with baby crying, that could count as hearing the same instance at two different times. Hero, need not, in this case, recognize any ordering to the Mater Sounds. There could, of course be a case in which he did, that is, in which the three Master Sounds were a saxophone playing a *C*, a *C#* and a *D*. But the point here is that Hero need not recognize any ordering of the Master Sound in order to identify and reidentify an instance of *Jesu* with respect to it.

Now one might object, in a way suggested by Evans (1985, p. 253), that such a move does not employ any real notion of numerical identity as distinct from a notion of qualitative identity, since if occasions of Jesu are distinguished in virtue of co-occurring with different Master Sounds (a baby cry and a saxophone blast), then the auditory experiences in question are qualitatively distinguishable. Thus reliance on the Master Sound seems not to be a reliance on genuine criteria of numerical identity.

As a response against Evans, I note that this problem arises only by treating the experience of the Master Sound at a given moment and the other, non-Master, sound heard at that moment to be blended into a single particular. Allowing that this need not be the case, that is, allowing that Hero can distinguish between the Master Sound at a moment and a non-Master-Sound at the same moment resurrects the possibility of the application of genuine criteria of numerical identity. Perhaps the Master Sound and non-master sounds are distinguished by timbre: the master sounds are a perfect sine-wave, and

a perfect saw-tooth wave, whereas the non-master sounds are saxophone performances. By treating the perfect sine wave and saw tooth wave tones and saxophone performances as different particulars, Hero can make sense of qualitatively identical but numerically distinct saxophone performances. This is not to say that Hero's experience is such that a concept of the numerical/qualitative distinction may be derived from it, but instead that it supplies materials out of which Hero may stipulate criteria for the employment of such a concept. One might object, as Grush does (personal communication) that such criteria nonetheless fail to provide for full blown numerical identity, since Hero would be distinguishing particulars based on perceptible qualitative differences. I take it that this objection amounts to asserting that numerical non-identity is compatible with total indistinguishability. My reply to this sort of objection is that such a notion of identity seems of little use in support of Strawson's thesis regarding space, since even spatial criteria are incapable of supporting the existence of absolutely indistinguishable non-identicals.

It seems, then, that Hero can apply criteria of particular reidentifiability without conceiving of the particulars as being ordered along any dimension. And insofar as dimensionality is necessary condition for spatiality, reidentifiability does not require space.

It is puzzling that Strawson originally stipulated that Hero's experiences were devoid of spatial content: while satisfying some of the necessary conditions of real space, Strawson at least would grant that not all of them were satisfied in Auditoria. It seems then that even if Strawson succeeded in showing that Hero did need to treat an ordered series of master sounds as criteria for the reidentifiability of song performances, this is insufficient to make the criteria spatial.

Before leaving this section I want to consider some further possible objections against me. It may be objected that I have, in discussing space and reidentifiability, overlooked a way that Strawson characterizes space and thus not done justice to his Reidentifiability Argument. Strawson and Evans both characterize space as that in virtue of which different things can simultaneously exhibit a system of relations over and above those which arise from the definite (intrinsic, non-relational) character of each (Strawson 1959, p. 79; Evans 1985, p. 253). They say little to flesh out this sparse characterization. Nonetheless, this characterization, with or without flesh, is insufficient to characterize real space. Below I show that this characterization can be satisfied by a system of non-spatial relations.

Strawson and Evans do not unpack the "over and above" characterization but I think that the idea is essentially the following. In Hero's theory of the world, he subscribes to some statements the predicates of which are monadic. To these correspond the intrinsic properties of objects. Other predicates in Hero's theory are binary: to these

correspond the relations. To unpack the notion of "relations over and above those which arise from the intrinsic character" of the relata, I propose the following. Let us suppose that in Hero's scheme, he has only the following kinds of predicates: monadic predicates for pitch and timbre, and the binary predicates "*x* has a higher pitch than *y*" and "*x* is louder than *y*". Hero's theory contains theorems regarding the "*x* has a higher pitch than *y*" predicate that make its application depend on the applicability of the monadic pitch predicates. For example, it may be a theorem for Hero that if *x* is a *C* and *y* is a *C#*, then *y* has a higher pitch than *x*. Thus, relative to Hero's conceptual scheme, the relation of one sound's having a higher pitch than another is a relation that arises out of the intrinsic characters of the sounds. In contrast, Hero has theorems regarding which sounds are louder than others, but the application of the relational predicate "*x* is louder than *y*" is not contingent on what pitch or timbre the sounds happen to be. Thus the relation of one sound's being louder than another, as conceived of in Hero's scheme, is my best guess as to what Strawson and Evans might mean by relations over and above those which arise from the definite (intrinsic, non-relational) character of each. This is not to say that I endorse this or any account of the distinction between intrinsic and extrinsic properties. But I propose to grant Evans and Strawson the distinction and focus on a different aspect of their account.

The next question to ask is the following. Can volume underwrite a set of relations between things (sounds) that arise over and above those that are due to the intrinsic natures of those things? Let us suppose that Hero conceives of the intrinsic

properties of sounds as those of pitch and timbre. Suppose also that Hero conceives of different sounds as bearing louder than relations to each other that are over and above those that arise from the intrinsic character of each. Contrast these relations to those such as *x* has a higher pitch than *y*, which, in Hero's conceptual scheme, do arise from the intrinsic character of each. Hero can do all of this without conceiving of the sounds as instantiating real spatial properties and relations, that is, without utilizing genuine spatial predicates in his theory of his world. Thus the "over and above" characterization is insufficient to distinguish real space from quasi-spaces and any argument that does not go beyond the "over and above" characterization of space is insufficient to show that real space is a requirement for objectivity.

Another line of objection, suggested by Grush (personal communication), advocates what I have called the "weaker" readings of the thrust of Strawson's arguments, especially the reidentifiability argument. One such reading would be that if some particular is objective, there must be some particular or other that is reidentifiable, and reidentifiability requires having spatial properties. Another reading would be that if some particular is objective, it must be reidentifiable, and reidentifiability requires that some particular or other has spatial properties. Yet another would be that if some particular or other is objective, then some particular or other is reidentifiable, and if some particular or other is reidentifiable then some particular or other is spatial. Are any of these weaker readings of the reidentifiability argument supported by the particular sections of text in which it appears? It seems not. For instance, throughout Strawson's

discussion of the first part of the reidentifiability argument, the part concerning whether objectivity requires reidentifiability, the question addressed within the context of Auditoria is whether the objectivity of particular sounds requires the reidentifiability of sounds—the very same particulars, not, as the weaker reading would have it, the reidentifiability of some other (non-audible) particular. Further, when Strawson takes himself to have provided for criteria of reidentifiability by providing the Master Sound, again, the very same particulars that are to be reidentified are the ones whose spatiality and objectivity are called into question. The actual structure of Strawson's arguments in chapter 2, especially the reidentifiability argument, conform to the strong reading, not the weaker readings suggested by remarks Strawson makes in chapter 1.

4.2 Strawson's Elsewhere Argument

In order for us to conceive of unperceivable particulars, we must, according to Strawson, have some understanding of why they might be unperceived. Before Strawson turns to consider how this necessary condition might be met in Auditoria, he asks how it is met in our familiar world. His answer, unsurprisingly, emphasizes the importance of the imperceptibility of the spatially distant, that is, the importance of

> a spatial system of objects, through which oneself, another object, moves, but which extends beyond the limits of one's observation at any moment.... Thus the most familiar and easily understood sense in which

> there exist sounds that I do not now hear is this: that there are places at which those sounds are audible, but these are places at which I am not now stationed. (1959, p. 74)

Strawson considers and dismisses the following non-spatial alternative. Perhaps a subject can think of unobservable existence as a function of the failure of sensory powers. Strawson dismisses this suggestion on the grounds that such an alternative cannot allow for being able to distinguish the failings of sense from the fading of the world. According to Strawson, the application of such a distinction requires spatial criteria. Unperceived particulars must be elsewhere.

I offer, contra Strawson, that we are able to conceive of objective things without conceiving of them as being elsewhere. I can conceive of myself as existing unperceived. I can conceive of myself as being knocked unconscious and locked in a darkened cellar. It is quite conceivable that several hours could pass without my being perceived by anyone, not even myself. And certainly I cannot help but be where I am. As Buckaroo Bonzai says: "Where ever you go, there you are". Thus there is at least one thing I can conceive of as existing unperceived without that thing being elsewhere: me. The Elsewhere Argument goes nowhere.

One may attempt to defend Strawson on this issue as Grush (personal communication) does by suggesting that the conceivability of an objective particular that does not exist elsewhere is a conceptual capacity that depends on the fact that I already have at my disposal spatial concepts. If it can be shown that conceiving of myself as able

to exist unperceived depends on having spatial concepts, then this would certainly help Strawson. But it must be noted that it has not been shown. Neither, of course, has its negation been shown. But something else can be said against this line of objection. This line of objection takes a form that is much too strong to be useful in defending Strawson's position. Strawson's arguments concerning objectivity take the form of showing that space is necessary because *such-and-such* is inconceivable. Putative counterexamples that *such-and-such* is conceivable are countered on this line of objection by asserting that *such-and-such* is conceivable only because we already have spatial concepts. But any situation in which we are to evaluate Strawson's arguments will be one in which we have spatial concepts. If it is admissible to doubt counterexamples to Strawson's arguments on the grounds that we already have spatial concepts, then it is likewise admissible to doubt Strawson's arguments themselves on the same grounds. Any claim about what may or may not be conceivable to a creature lacking spatial concepts may not be adequately evaluated by creatures such as us that already have spatial concepts. One begins to wonder, then, what the point of Strawson's exercise is. If it is granted that we already have spatial concepts, how could it be shown that they are necessary for certain things if we cannot meaningfully contemplate lacking spatial concepts? Strawson's own methodology demands that counterexamples to his claims be taken seriously in spite of being generated by creatures that already have spatial concepts. Barring independent argument that *such-and-such* is conceivable only because we have spatial concepts, the bare suggestion should not give us further pause.

4.3 Evans' Simultaneity Argument

According to Evans, an essential feature of our concept of an objective world of spatially located objects is that of many things, perceived and unperceived, existing simultaneously. This is integral to objectivity insofar as conceiving of existence unperceived involves conceiving of the objects unperceived as existing at the same time that they are being thought about though not perceived. Simultaneity may be a concept whose application is clearest in visual experience. Vision allows us to be simultaneously aware of distinct objects. Such experiences afford an opportunity for the direct application of what Evans calls "simultaneous spatial concepts" (1985b, p. 283). Simultaneous spatial concepts are concepts of the spatial relations between two or more relata the direct application of which requires the simultaneous perceptual experience of all of the relata. In contrast, "serial spatial concepts" require only the non-simultaneous experience of the spatial relata for their direct application. For examples of the latter kinds of experiences, Evans asks us to consider the way a blind person may come to know the spatial configuration of a large object such as a table by running his hands over its edges and surfaces. Here the experience of all of the table's features are not perceived

simultaneously but instead, in succession.[4] Such kinds of experience seldom afford opportunities for the direct application of simultaneous spatial concepts. Relatively small objects, however, can be taken in tactually all at once, as when a small cube is cupped in the hands. According to Evans, a subject of purely auditory experiences, however, never has the opportunity for the direct application of simultaneity concepts and thus could never have simultaneity concepts. Thus, according to Evans, such a subject is barred from having genuinely objective experiences, since objectivity entails simultaneity and Hero's experience of his auditory quasi-space does not allow for genuinely simultaneous experiences.

I propose to resist Evans' denial of Hero's application of simultaneity concepts as follows. I grant that simultaneity is a feature of real space. However, the mere fact that some set of properties may be instantiated simultaneously is insufficient to make those properties spatial. A purple object may be regarded as the simultaneous instantiation of two different things: a red one and a blue one. Consider an ontology of volumes of colored gasses in which red and blue volumes are basic—not created by the mixing of any other volumes—but what we would perceive as purple volumes appear only upon the intermingling of equal parts of red and blue. On such a theory of objects it would be

---

[4] Evans' point concerning the blind is not that the tactile experience of the parts of a large object never afford the application of simultaneity concepts, but instead that they do not afford the *direct application* of such concepts.

legitimate to allow for two different objects, one red and one blue, to occupy all and only the same spatial locations, as opposed to admitting a third kind of object, a purple one, into the ontology. Again it must be stressed that this is not a claim about anything but admissible judgment-calls that may be made, a point dealt with above. The co-perception of a red object and a blue object need not be a perception of the objects as being located at two different places. To make the same point in terms of the "sounds" thought experiment, consider the following.

Suppose that, until today, Hero has been listening to a series of alternating HONKs and DINGs. Imagine that the HONKs are bass blasts from a baritone saxophone and that the DINGs are tinny tones from a diminutive xylophone. Hero is hearing the following series: HONK DING HONK DING HONK DING. Suppose that one day Hero hears a sound that is qualitatively identical to what we would hear if a DING and a HONK occurred simultaneously. Now it seems that Hero could interpret this occurrence in one of two ways, one of which involves the application of a concept of simultaneity, the other does not:

> *The Simultaneity Option*: I (Hero) just heard two different features simultaneously: a HONK and a DING
>
> *The Non-Simultaneity Option*: I (Hero) just heard a (third) feature that I've never heard before: a DONK.

Now, according to Evans, Hero is barred from the simultaneity option. But this invites the following question: Why should Hero conceive of today's experience as a DONK

rather than as a simultaneous occurrence of a HONK and a DING? Evans' take on this issue has not been argued for, nor does it seem obviously correct.

I am prepared to grant that if Hero only heard HONKs and DINGs simultaneously, he would never be compelled to interpret them as anything other than the presentation of the single feature DONK. However, given the supposition that he frequently, if not usually, hears them separately, then I see no reason for him to be barred from hypothesizing that DONK-ish experiences are actually the simultaneous occurrence of HONKs and DINGs, i.e., that DONKs reduce to HONKs and DINGs. It seems, then, that Hero can conceive of distinct events occurring simultaneously without having to conceive of them in different spatial locations. Thus, even if objectivity required simultaneity, as Evans suggests, simultaneity does not require space. Thus the road from objectivity through simultaneity to space is blocked at the path from simultaneity to space.

Grush suggests a line of objection to my treatment of the Simultaneity Argument. The gist of the complaint is that I misinterpret Evans' core claim. According to this line, Evans is not claiming that Hero couldn't have simultaneity concepts, nor is Evans claiming that Hero would be incapable of directly applying simultaneity concepts to experience. Instead, the point is that "nothing in the description of Hero's experience requires us to think of him as having such simultaneity concepts, and without them, his experience could not be objective" (Grush, personal communication). However, textual

evidence does not support this line. The key passage concerns Evans comparing blind humans to Hero. The blind, while lacking opportunity for the direct application of simultaneity concepts, may nonetheless be situated to apply them indirectly. Describing the blind, Evans writes that

> it must be supposed that they are equipped with the appropriate concepts and that they 'interpret' or 'synthesize' their sequential experience in terms of them. This is independently plausible in the case of the blind; but appears to have no independent plausibility in the case of the subject of a purely auditory experience, who similarly, and indeed far more certainly, lacks any opportunity for direct application of simultaneous spatial concepts." (Evans 1985b, p. 276).

Here it is quite clear that Evans' aim is to deny Hero any opportunity for the application of the requisite concepts to experience. Evans' is not merely denying that experience could constitute a source of such concepts, a font from which such concepts necessarily flow.

4.4 Evans' Causal Ground Argument.

This argument represents a departure from the previous three in that Evans attention is turned from the metaphysical objectivity of objects (i.e. particulars) to the metaphysical objectivity of properties.

Evans argues that the properties we perceive objects as having—properties Evans calls "sensory properties"—can be conceived as instantiated unperceived only if they are conceived as causally grounded in spatial properties.

Evans' argument can be seen as having two parts. The first part is the suggestion that we cannot conceive of the sensory properties as being instantiated unperceived without supposing that they instantiate these properties in virtue of instantiating some other properties. The second part is the suggestion that these other properties—the causal ground of the sensory properties—must be spatial properties.

The two parts open the argument to two lines of attack. The first is to question the necessity of supposing unperceived sensory properties being co-instantiated with some other properties. I pursue this line of attack at some length below. The second line of attack questions the need for these other properties to be spatial. Even if Evans is correct that sensory properties must be conceived of as having a causal ground, why must this causal ground be thought of as consisting of spatial properties? I do not pursue this second line of attack beyond pointing out here that Evans seems not to have argued for the necessity of the spatiality of the causal ground. I am more interested in pursuing a third line of attack, namely, to point out an incoherence in the way Evans' describes the objectivity of sensory properties. In the remainder of this section I offer two arguments that flesh out the first and third lines of attack, respectively. First, I argue that the objectivity of so-called sensory properties does not require them to have a causal ground.

Second, I argue that Evans' characterization of sensory properties renders incoherent the suggestion that we conceive of them as objective properties.

Before turning to my arguments against Evans, I want to remark upon the relevance of the Causal Ground argument to

> SO2. A property is metaphysically objective only if it is a spatial property.

At first glance it may seem that the Causal Argument is not relevant to SO2 since Evans' aim is to show that sensory properties are objective only if co-instantiated with spatial properties. To be relevant to SO2, the objectivity of sensory properties must require more than co-instantiation with spatial properties but instead identification with spatial properties. I think that Evans does allow for the eventual identification of the spatial causal ground of sensory properties and the sensory properties. Evans describes sensory properties (which he identifies with secondary properties) as follows.

> For an object to have such a property is for it to be such that, if certain sensitive beings were suitably situated, they would be affected with certain experiences, though this property may, in its turn, be identified with what we should normally regard as the ground of the disposition. (1985b, 268-269)

The main question in need of consideration now is whether properties as we perceive them can be instantiated unperceived. Evans argues that since sensory properties are secondary properties, there is a theoretical difficulty in imagining them instantiated unperceived. According to Evans the closest that we can come to imagining sensory properties instantiated unperceived is by imagining their non-sensory causal ground

unperceived. And further, this causal ground must be comprised of spatial properties. What I want to do in this section is block the very first move that Evans makes in this argument, that is, block the move that sensory properties by themselves cannot be imagined to be instantiated unperceived.

Evans argues that it is quite difficult to see how an object "as we see it" can be the same as when we do not see it (Evans 1985b, p. 272-274). Suppose that I am seeing an apple as red. How can it be red when no one is seeing it, when, say, it is locked in a dark cellar? Evans contends that this is inconceivable (Evans 1985b, p. 274). Evans writes that "All it can amount to for something to be red is that it be such that, if looked at in the normal conditions, it will appear red" (1985b, p. 272). Evans contrasts this view with one that tries "to make sense of the idea of a property of redness which is both an abiding property of the object, both perceived and unperceived, and yet 'exactly as we experience redness to be'"(1985b, p. 272). Evans objects against this latter view that "it would be quite obscure how a 'colour-as-we-see-it' can exist when we cannot see it, and how our experiences of colour would enable us to form a conception of such a state of affairs" (1985b, p. 273).

I want to defend this latter view by suggesting that the obscurity alleged by Evans arises due to a concealed ambiguity in sentences employing phrases like "as I see it". Once such phrases are properly disambiguated, it becomes quite clear how a color as we see it may be the same when it is not seen.

I begin by considering sentences employing phrases with the form "x as I am F-ing it". Consider the sentence

> The chair as I am standing next to is the same as when I am not standing next to it.

There is a reading of this sentence whereby it is quite clearly contradictory. On such a reading the sentence expresses the claim that a chair stood next to is a chair not stood next to. This is contradictory on the supposition that a chair cannot be both stood next to and not stood next to at the same time. Suppose, then, that we were to read the following sentences along similar lines.

> The chair as I see it is the same as when I do not see it.

On such a reading, Evans would be correct that it is quite obscure how the chair as I see it can be the same as when I do not see it. There is a difference between the chair as I see it and the chair when it is not seen by me, namely, in the first case I am seeing it and in the second I am not. And on the supposition that the chair cannot be both seen and unseen at the same time, the sentence under consideration expresses a contradiction.

Sentences employing phrases with the form "*x* as I am *F*-ing it" may be read in different way than considered so far. This alternative reading shows the non-contradiction of

> The chair as I see it is the same as when I do not see it.

I call your attention to an analogy between the above sentence and the following.

> The chair as I am describing it is the same as when I am not describing it.

This sentence admits of a reading whereby it expresses a contradiction. On such a reading the above sentence is equivalent to

> The chair described is not described.

But on the alternative reading I want to consider, a chair can be as I describe it even when I am not describing it.

Suppose that I am describing the chair as having been manufactured in Switzerland. I am uttering the sentence "This chair was manufactured in Switzerland". My describing the chair is just my uttering a sentence. The chairs' being as I describe it however, is not its being a chair in the proximity of someone uttering a sentence. The chair's being as I describe is, in this case, its having been manufactured in Switzerland. Clearly a chair may have been manufactured in Switzerland regardless of whether I am now describing it as such. With this last point in mind, then, we may read

> The chair as I am describing it is the same as when I am not describing it.

as non-contradictory on the grounds that a Swiss chair doesn't stop being Swiss when I stop talking.

To sum up what I have said so far, I would say that

> The chair as I am describing it is the same as when I am not describing it.

admits of two readings: a representational reading and a non-representational reading. On the representational reading, the sentence is not a contradiction. On the non-representational reading, the sentence is a contradiction. On the representational reading, I am predicating being manufactured in Switzerland of the chair. The chair may very well instantiate the property of being manufactured in Switzerland even when I am not representing it as such. On the non-representational reading,

> The chair as I am describing it is the same as when I am not describing it.

is analogous to

> The chair as I am standing next to it is the same as when I am not standing next to it.

and both are contradictory on the following grounds. A chair stood next to is different from a chair not stood next to and a chair described is different from a chair not described.

Evans detects obscurity and unintelligibility in the supposition that "a 'colour-as-we-see-it' can exist when we cannot see it" (1985b, p. 273). I suggest that the supposition is clear and intelligible if read representationally. If we explicate perception as a representational affair, then

> The chair as I see it is the same as when I am not seeing it

is no more a contradiction than the representational reading of

> The chair as I am describing it is the same when I am not describing it.

I may describe a chair as being Swiss and it can continue to be Swiss during periods when it is not described. Likewise, I can see a chair as being brown and it can continue to be brown during periods when it is not seen.

Allowing representational readings of sentences employing phrases of the form "*x* as I am *F*-ing it" shows how the first part of Evans' Causal Ground Argument is blocked. The first part of Evans' argument is his allegation that there is something obscure and unintelligible in the claim that objects can be the way that we perceive them even when unperceived. A representational analysis of perception renders the claim clear and intelligible. Thus the first move in Evans' Causal Ground Argument—the move by which Evans alleges to show that we cannot conceive of sensory properties by themselves being instantiated unperceived—is blocked by adopting a representational explication of perception.

Perhaps it may be thought that the representational explication is unavailable in precisely the situations where the properties in question are secondary (subjective), that perception can be representational only insofar as what are perceived are primary (objective) properties. I take it that one of the main burdens of the first three chapters of this dissertation is to show that this is not the case, that what may be a content of mental representations, including perceptual representations, may be metaphysically subjective in precisely the ways that secondary properties are traditionally thought of as

metaphysically subjective: as having their existence or defining essence depend on the representing subject.

What room is there for Evans to defend the Causal Ground argument by denying, as Grush (personal communication) does, the truth of the representational account of perception? There is little, if any room, for Evans to make such a maneuver. Note first that the logic of the situation is something like this. Evans claims, not possibly P and I counter, possibly P because representationalism is possibly true. Now, someone may attempt to save the Evansian line by denying the possible truth of representationalism. But Evans would not. The whole project of providing for the objectivity of perception presupposes that perceptual states are representational states: The notion of there being things that we perceive that may exist independently of our perceiving them is asked within representation milleu, within a context that presupposes that perceptual states are representational states. Further, evidence may be gathered showing that Evans actually does buy into representationalism about experience. In Evans 1982, perception has representational content, though Evans famously allows that in some cases that content may be nonconceptual and illustrates the case in terms of colors. Evans famously asks "Do we really understand the proposal that we have as many colour concepts as there are shades of colour that we can sensibly discriminate?" (1982, p. 229) thus launching the cottage industry of non-conceptual content studies. Now one may want to suggest, upon comparing this remark to the causal ground argument, that Evans' didn't think there was non-conceptual content after all, in spite of being historically credited with introducing

the notion into contemporary discussions. Pursuing such a grand project goes way beyond my purposes at hand, namely, to show that the causal ground argument does not force upon the predicational theory of objectivity any destructive consequences.[5] I turn now to my second problem with Evans' Causal Ground argument.

One of the upshots of a representational explication of perception is the allowance for a distinction between the properties of perceptions and the properties of the things perceived. This distinction is analogous to the distinction between my stating that water is wet and water's being wet. The supposition that a property is objective includes the supposition that its instantiation does not depend on its being represented. Why, then, does Evans think that there is a special problem in conceiving sensory properties as being objective–the properties we perceive objects as instantiating? The answer to this question, I suggest, is that the problem arises because Evans thinks that sensory properties are subjective properties. He describes them at length as secondary properties—properties whose defining essence is their disposition to cause certain experiences (1985b, pp. 268-272). On my view of metaphysical objectivity and subjectivity, secondary properties are

---

[5] It may be worth mentioning as an aside that "*x* is as I *F* it even when I'm not *F*-ing it" can be given sense apart from representationalism. Say I chop a stone in two. The way I've chopped it (in two) is a state of affairs that may have obtained without (that is, independent of) my chopping, it may have been struck by lightning or simply fell apart into two pieces.

prototypical instances of the metaphysically subjective. Thus the problem that Evans sets for himself in his Causal Ground argument is the problem of trying to figure out how it is that we conceive of subjective properties as being objective. But if we are prepared to admit that they are subjective, then it seems that we have lost interest in trying to conceive of them as objective. If sensory properties are not really objective after all, then they are an entirely useless platform from which to launch a defense of the spatial criterion of objectivity.

I should note that there is something unfair about my complaint as waged against Evans. Evans analyses colors as depending on a disposition to be experienced, but does not equate them with a property that is instantiated only when experienced. Evans' notion of objectivity in the Causal Ground Argument is no more than that of a property that can be instantiated even when unperceived. Both Evans and I allow that secondary properties may be instantiated unperceived. But unlike me, Evans does not regard secondary properties as subjective, but instead, objective. In fairness, the remarks in the current section should not be regarded so much as critical of Evans, but instead as a way of showing that, as I use the terms "objective" and "subjective", secondary properties cannot serve as a platform upon which to erect an argument for the spatial criteria of objectivity.

I have argued for the failure Evans' and Strawson's arguments that the objectivity of objects (particulars) and properties requires them to be spatial. Evans and Strawson thus do not provide compelling reasons for amending the predicational theory of

objectivity to include spatial requirements on the metaphysical objectivity of objects (SO3) and properties (SO2). If Evans' Causal Ground argument for the spatial requirements of the metaphysical objectivity of properties had succeeded, then it would have served—given the predicational theory—as an argument that epistemically objective intentional phenomena must have predicates that name spatial properties (SO1). Evans and Strawson's arguments fail to provide compelling reasons for the inclusion of SO1-SO3 in the predicational theory of epistemic objectivity.

Before leaving the discussion of Strawson and Evans, I want to consider one final objection to the line of argument in this chapter, namely, that the notion of objectivity that I have been primarily concerned with is orthogonal to that which is the focus of Strawson and Evans' arguments (Grush, personal communication). What is the alleged contrast? Strawson writes:

> The limit I want to impose on my general question is this: that I intend it as a question about the conditions of the possibility of identifying thought about particulars distinguished by the thinker from himself and from his own experiences or states of mind, and regarded as actual or possible *objects* of those experiences. I shall henceforth use the phrase, 'objective particulars' as an abbreviation of the entire phrase, 'particulars distinguished by the thinker &c.' (1959, p. 60).

Thus, Strawson is interested how it is that we conceive of particulars existing unperceived. Is this different than what I've been interested in all along? Grush (personal communication) urges that the contrast between Strawson's target and my own is that I am interested in particulars that exist independently of perception and Strawson is interested in particulars *conceived* to exist independently of perception. I think, however,

that Strawson and I are not talking past each other. Recall that the project I set for myself in chapter 1 included explicating the concept of objectivity, which involves unpacking what it is to conceive of things as existing unperceived, or more broadly, existing mind-independently. Strawson and I may diverge over what the phrase "objective particular" is shorthand for, but ulitmately we are interested in the same thing. We are interested in what it means, *relative to the conceptual scheme of humans,* for there to be a distinction between our experiences and the mind-independent objects of experience.

One might urge an even stronger claim here, one that even Strawson would be prepared to accept, namely that there really is no difference between "exists independent of us" and "relative to our conceptual scheme, exists independent of us". All claims that we can evaluate or even consider are claims relative to our conceptual scheme. Does such a concession threaten realism, or any hope for genuine objectivity? Such will be the focus of the next chapter.

Chapter 5 The Epistemology of Objectivity

5.0 Introduction

We are natural objects in a natural world. We notice things, and we notice ourselves among those things. And, we notice our own noticings. Our ways of taking in the world are themselves taken in. Among the objects of our scrutiny are our own scrutinizings. Of the things in the world, there are those things that are us, and those things that are not us. Among the things that are not us are those things that nonetheless depend on us. These things are subjective, and all else, by contrast, is objective. My main concerns in previous chapters were metaphysical: what is it to be objective or subjective? I turn now to epistemological questions, in particular: How do we know which things are objective and which are subjective?

In this chapter I explore answers to such epistemological questions within a naturalized epistemology—that is, within a framework that sees the first, last, and best answers to "how do you know. . .?" questions coming from—or at least being cast in the idiom of—the empirical sciences.

One of my key motivations for writing the current chapter is to ward off a kind of response that I have encountered often in discussing earlier versions of the above material. This general kind of response sees my project as attempting (and failing) to defend classical philosophical positions, like realism or anti-skepticism. While there is much continuity between the language and issues of the current work and these classical debates, taking up these classical causes it outside the limit of what the naturalist may accomplish or even care about. My goal instead is to show, relative to a scientific viewpoint, the utility of treating some things as mind-dependent (in the idealist sense) and some things and mind-independent (in the realist sense). The questions to be addressed here are: assuming a scientific view how would you know what is real and what is ideal?

5.1 The Logic of realism and idealism

In this section I unpack the notions of mind-dependence and independence by looking at the ways such notions are treated in debates between realists and anti-realists. Focusing on the key logical structures of such debates will allow for clarity regarding the issues to be pursued in the sections that follow. Debates over realism are typically local, not global. That is, in a typical realist/anti-realist debate, the realism in question is restricted to, say, moral realism, or scientific realism or realism about mathematical objects, to name but a few local battlefields. Thus, typically, one is not a realist

simpliciter but a realist about entities in some restricted domain of discourse. One may find someone who is a realist about electrons, but an anti-realist about aesthetic properties. Despite the varieties of realism and antirealism correlative with the varieties of domains of discourse, there is a common logic to realist/anti-realist debates. This general logic may be illustrated with the following example. Being a realist about electrons involves committing to a two part claim: (i) that electrons exist and (ii) that electrons exist mind-independently that is, independent of whether anyone talks, thinks, or theorizes about electrons. Given that the realist thesis is twofold, anti-realists are of two kinds.

The first kind of antirealist denies (i). For example, they deny that that there are any electrons. This kind of antirealist is a nihilist. Hartry Field is a nihilist about numbers and Bas van Fraassen is a nihilist about electrons (Field, 1989; van Fraassen, 1980). Quine is neither a nihilist about numbers nor electrons (Quine, 1980). This is insufficient, however, to make him a realist about them, since the question about mind-independence remains undressed.

The second kind of antirealist denies (ii). Thus, whether there are horses, for instance, is said to depend on whether anyone talks, thinks, or theorizes about horses. The second kind of antirealist is an idealist. George Berkely, Immanuel Kant, Nelson Goodman, Richard Rorty, and Hilary Putnam are all idealists. Berekely famously remarked that to be is to be perceived. Thus whether horses exist depends on whether

horses are experienced by some mind or another. Equally famously, Quine remarks that to be is to be the value of a bound variable. Thus he seems to assert that whether horses exist depends on whether horses are quantified over in some theory or another. Thus Quine's ontological relativity is a kind of idealism insofar what exists is relative to what is said to exist. These are of course matters of interpretation. What matters is not the correctness of the interpretations, but the illustrative purposes they serve.

For the sake of simplicity, I will discuss realism and idealism about electrons, although I intend my remarks to apply to realist/idealist debates about, say, horses or numbers as well. Also for simplicity's sake, I will express the idealist thesis in terms of the mind-dependence of electrons. Various versions of idealism about electrons will differ over what electrons are dependent on—theories, languages, perceptions, thoughts, and cultures are just a few of the considered options. I use 'mind' as shorthand for these options. In this section I am much more concerned with the notion of dependence at play in these debates than the notion of mind.

Given the above conventions of shorthand I adopt, then, the disagreement between idealists and realists is over whether the existence of electrons depends on the existence of minds. My focus here is on what this dependence amounts to. I argue that this dependence exceeds anything expressible in a solely extensionalist idiom. That is, the dependence of electrons on minds asserted or denied by parties to the debate cannot be stated merely in terms of material (truth functional) dependence. Thus, the

investigation of the real and ideal will have to avail itself of means of expression that outstrip extensional language, and utilize instead, the language of causal and counterfactual dependence, topics best explored in a naturalistic mind-set and taken up in the next section. If the realists and idealists disagree, then the idealist claim cannot simply be the material conditional statement that if there are electrons then there are minds. Nor can the realist's claim simply be the denial of the truth of this material conditional. Let us begin by supposing it were otherwise. Suppose, say, that the idealist tried to express her thesis using

(1) $(\exists x)(Ex) \rightarrow (\exists y)(My)$

as an extensional expression of his thesis where '*Ex*' is '*x* is an electron', '*Mx*' is '*x* is a mind', and '→' is the truth functional connective of the material conditional. We might read (1) as 'There is something that is an electron only if there is something that is a mind'. Analogously, we might characterize the realist's thesis as

(2) $\sim[(\exists x)(Ex) \rightarrow (\exists y)(My)]$

where '~' is the truth functional negation operator.

A problem arises, however: (1) and (2) are inadequate to capture the disagreement between the realist and the idealist. This is because the realist as well as the idealist assents to the truth of (1). To see why, consider that material conditionals are false only

if their antecedents are true and their consequents are false. But the realist and idealist each grant the truth of the antecedent and consequent of (1). They agree that, as a matter of fact, it so happens that there are electrons and minds.

What the realist and idealist disagree about goes beyond the mere coexistence of electrons and minds. The idealist insists that in order for there to be electrons there must be minds. According to realist, on the other hand, there could have been electrons even if there were no minds. As I will argue in the sections that follow, the idealist and the realist rely on modal notions to express their respective theses. The idealist asserts

(3) $[\Box(\exists x)(Ex) \rightarrow (\exists y)(My)]$

which reads 'Necessarily, if there is something that is an electron, then there is something that is a mind'. The realist asserts

(4) $\sim [(\Box\exists x)(Ex) \rightarrow (\Box\exists y)(My)]$,

which is equivalent to

(5) $\sim[(\Box\exists x)(Ex) \rightarrow (\Box\exists y)(My)]$

and

(6) $[(\exists x)(Ex) \& \sim(\exists y)(My)]$.

The realist, then, believes it possible that there be electrons without minds where the idealist holds that if there are electrons, then that there are no minds is impossible. One of the clearest modern idealist arguments is due to George Berkeley (1710, 1713). Of Berkeley's arguments for idealism, the one that he regarded as central to his system[6], and the one that has come to be called his 'master argument' (Pappas 1995 p. 73) is the following. If one can "conceive it possible" that the objects one thinks about "may exist without the mind", then "it is necessary that [one] conceive them existing unconceived of or unthought of, which is a manifest repugnancy" (Berkeley 1710, §23). The argument is echoed in the persona of Philonous, Berkeley's spokesman against the realist Hylas in The three dialogues.

> *Hyl*. . . .What is more easy than to conceive a tree or house existing by itself, independent of. . .any mind whatsoever? I do at this present time conceive them existing after that matter.
>
> *Phil*. How say you, Hylas, can you see a thing which is at the same time unseen?
>
> *Hyl*. No, that were a contradiction.

---

[6] While Berkeley has other arguments for his idealism (e.g. the argument from perceptual relativity) he is "content to put the whole upon this issue" (1710, § 22).

Phil. Is it not as great a contradiction to talk of conceiving a thing which is unconceived?

Hyl. It is.

Phil. The tree or house therefore which you think of, is conceived by you.

Hyl. How should it be otherwise?

[. . .]

Phil. How then came you to say, you conceived a house or tree existing independent. . . of all minds whatsoever?

Hyl. That was I own an oversight. . . . (1713, First Dialogue)

The Berkeleyean theme exemplified in his master argument is also echoed in contemporary idealist remarks. Hilary Putnam writes:

> I agree with Rorty that we have no access to "unconceptualized reality." As John McDowell likes to put it, you can't view your language "from sideways on" in the way that the idea of looking at one's language and the world and comparing the two suggests. (1994, p. 297)

According to worldmaker Nelson Goodman, "We cannot test a version by comparing it with a world undescribed, undepicted, unpercieved. . ." (1978, p. 4). Elsewhere, describing his theory of 'conceptual relativity,' Putnam writes ,

> Our language cannot be divided up into two parts, a part that describes the world "as it is anyway" and a part that describes our conceptual contribution. This does not mean that reality is hidden or noumenal: it simply means that you can't describe the world without describing it. (1992, p. 123)

Everything that we think about is thereby thought about. Who could deny that? Everything that exists is thought about. Thinking the thought expressed by the previous sentence makes it the case that everything that exists is thereby thought about. Are these remarks sufficient to establish idealism? Are they sufficient to establish that what exists depends for its existence on being thought about? No. All they establish is the material equivalence of that which exists and thoughts about it, which is something already granted by the realist.

What, then, can the realist respond to the idealist? The realist cannot give an example of anything that is not thought about, because as Philonous would be quick to point out, any attempt would be self refuting.

The first step of any realist response will be to point out that at best the idealist has shown that any electrons we have thought of have been thought of. And this is a matter of fact with which the realist has no quibble. The realist thesis goes beyond this simple fact. The realist will grant that the only reality that we conceive of or talk about is thereby a reality conceived of or talked about. But, the realist will insist, this is insufficient to make it the case that reality *must* be conceived of or talked about.

The realist has only three options for expressing his thesis, as far as I can tell. I call these three options the essentialist option, the counterfactual option, and the possible worlds option. And for each of these three options, the idealist has a corresponding way of expressing his opposition to the realist. I discuss these three pairs of options below.

The Essentialist Option

On this option the realist says that even though electrons cohabit the universe with minds, that is only a contingent or accidental attribute of electrons: cohabiting the universe with minds counts among none of the essential attributes of electrons. Thus the realist asserts

(7) Of the essential attributes of electrons (e.g, having a negative charge), cohabiting the universe with minds is not one of them.

And the idealist counters with

(8) That there are minds is an essential attribute of electrons.

The Possible Worlds Option

On this option, the realist and idealist cast their disagreement in terms of how many possible worlds that have electrons also have minds. Thus, according to the realist

(9) There is at least one possible world with electrons but no minds.

And according to the idealist

(10) Every possible world that has electrons also has minds.

The Counterfactual Option

On this option, the realist and idealist assert contrary counterfactual conditional about minds and electrons.[7] According to the realist, then,

(11) Electrons would have existed even if no minds had existed.

On the other hand, the idealist asserts

(12) If no minds had existed, then no electrons would have existed.

The three options (counterfactual, essentialist, and possible worlds) involve a notion of dependence stronger than mere material dependence. For simplicity's sake, I will focus on the counterfactual dependence: a property is counterfactually mind-dependent just in case its instantiation is counterfactually dependent on being represented. An example of such a property is the property of being seen. A rock at the bottom of the ocean cannot instantiate the property of being seen unless it is represented. The rock's being seen is counterfactually dependent on the rock's being represented. We can't help but be idealists about the seen things. Might there be properties of rocks that we are realists about? Are there any properties of rocks that are not counterfactually

---

[7] See Schmitt (1995, pp. 40-46) (a realist) for a discussion of the debate between realists and idealists in terms of disagreements about counterfactuals. (Schmitt's example is "Trilobites would have existed even if no minds had existed.")

dependent on minds? Arguably, whether a rock is made of basalt is not counterfactually dependent on minds. On our best accounts of what basalt is and what minds are, we must allow that there were rocks made of basalt long before there were any minds (at least, so says the basalt-realist).

Counterfactual mind-dependence is not the only kind of mind-dependence that has been glossed in the literature. There is another kind of mind-dependence, what I shall call weak mind-dependence (contrasting with counter-factual mind-dependence, which I shall call strong mind-dependence). Weak mind-dependence may be introduced by considering a property like being visible. Being visible is not counterfactually dependent on being seen, since, certain things were visible before they were seen. But we can understand what it is for something to be visible only by reference to a certain class of perceivers. In other words, if there is nothing more to having some property than having some disposition to elicit certain kind of response from a certain kind of mind, then is mind-dependent in this latter, weaker, sense.

The distinction between strong and weak mind dependence is similar to one that Dennett (1991) discusses. Dennett writes of two kinds of mind dependent properties—lovely and suspect properties. These kinds differ in that suspect properties are more strongly dependent on minds than lovely properties. Suspect properties, like the property of being a suspect, cannot be instantiated by some entity unless that entity actually is represented. Someone could not be a suspect unless they were actually

suspected of something. Likewise for the property of being seen. Nothing is seen unless it is visually perceived, and thus represented. In contrast, lovely properties are defined in terms of their dispositions to elicit mental responses. Someone may be lovely even though they lived in isolation their entire life and were never perceived to be lovely. Dispositional properties may be instantiated without being actualized: that is, something can be breakable without being broken. Dispositions defined in terms of representations thus do not need to actually be represented to be instantiated. Someone's being lovely is defined in terms of dispositions to elicit a certain response in a certain class of perceivers. Where the property of being seen is suspect property, the property of being visible is a lovely property. Dennett writes:

> Particular instances of lovely qualities (such as the quality of loveliness) can be said to exist as Lockean dispositions prior to the moment (if any) where they exercise their power over an observer, producing the defining effect therein. Thus some unseen woman (self-raised on a desert island, I guess) could be genuinely lovely, having the dispositional power to affect normal observers of a certain class in a certain way, in spite of never having the opportunity to do so. But lovely qualities cannot be defined independently of the proclivities, susceptibilities, or dispositions of a class of observers. Actually, that is a bit too strong. Lovely qualities would not be defined—there would be no point in defining them, in contrast to all the other logically possible gerrymandered properties—independently of such a class of observers. So while it might be logically possible ("in retrospect" one might say) to gather color property instances together by something like brute force enumeration, the reasons for singling out such properties (for instance, in order to explain certain causal regularities in a set of curiously complicated objects) depend on the existence of the class of observers. [. . .]
> Like Lockean secondary qualities in general, they are equivalence classes of complexes of primary qualities of those states, and thus can exist independently of any observer, but since the equivalence classes of different complexes that compose the property are gathered by their characteristic effect on normal observers, it makes no sense to single

> them out as properties in the absence of the class of observers. There wouldn't be colors at all if there weren't observers with color vision, and there wouldn't be pains at all if there weren't subjects capable of conscious experience of pains, but that does not make either colors or pains into suspect properties. (1991, Cited text from electronic version available on Dennett's web site.)

The mere fact that something has a disposition to be represented is insufficient to render it even weakly-mind dependent. Samples of gold may be such that if something is gold, then it thereby has a disposition to be thought of as such. This alone is does not make being gold a lovely (weakly mind-dependent) property. However, it would be a lovely property only if there was nothing *more* to being gold than having the disposition to be thought of as gold.

I want to redescribe metaphysical and epistemic objectivity in such a way that will allow me to broach certain topics about the epistemology of objectivity. It will be useful to restate the notions of strong and weak mind dependence in terms of the kinds and numbers of causal interactions that different properties may enter into.

In keeping with the causal covariational account of representation discussed in chapter 2, will say that a representation R represents a property P only if P is able to enter into causal interaction with R. With this notion in mind, let us redescribe the notions of epistemic and metaphysical objectivity/subjectivity, then. A metaphysically subjective property is a property that enters into causal interactions only with mental representations. (A maximally subjective property would be a property that enters into

causal interactions only with the predicate that expresses it.) An objective property is one that enters into causal interactions other than interactions with representations.

It may be worth while mentioning the kinds of causal interactions that may relate a predicative representation and the property it denotes. (Something very like the following catalog may be found in Fodor (1998, pp. 78-80)). The first and most noteworthy is *perception*. If one is able to mentally represent dogs, that is, predicate *is a dog* of some individual, then one is typically (perhaps necessarily) able to apply the representation in perceptual situations, in situations in which one is able to perceive the presence of a dog. One may, in a certain conditions, if presented with a dog, be able to perceive that it is a dog. Many different philosophers engaged in many different projects may wonder what conditions one has to be in. But, as will eventually become apparent, that does not matter here. What does matter is the truth of the following widespread assumptions: (a) Perceiving a dog involves representing a dog and (b) perceiving a dog involves being caused by a dog to token a representation of a dog. Precisely what kinds of causal mechanisms relate percepts to perceived things, I do not know. What matters here is only that there are some. It is worth pausing to preview how this might relate to some topics to be covered later, topics like publicity and objectivity. Consider the causal mechanisms that perceptually relate Helen Keller's dog-representation to dogs and another set of mechanisms that perceptually relate my dog-representations to dogs. In my case, the typical conditions in which I find myself perceptually representing dogs will be those conducive to vision: good lighting, etc. Such conditions will be quite irrelevant to

Helen Keller. For her, the conditions will involve those conducive to smelling and touching dogs. People with very different sensory abilities, myself and Helen Keller, for example, can nonetheless perceive dogs as dogs: there are multiple kinds of conditions in which the presence of a dog can causally elicit a representation of a dog. This is because there are multiple kinds of causal interactions that relate the property of being a dog to dog-representations. *A fortiori*, then, there are multiple kinds of causal interactions that the property of being a dog enters into. This is a necessary condition on a property's being objective, and being public. This will be explored in more detail, but I mention it now to give the flavor of where this is going. Entering into multiple causal interactions is a road to objectivity. And entering into multiple interactions relating a thing to representations of it is the road to publicity. I return now to cataloging the many kinds of causal interactions that may relate representeds to representations.

Perceptual mechanisms do not supply the only kinds of causal chains that may link a dog to my coming to token a dog-representation. Among the other kinds of linking mechanisms, I will briefly mention three more: Theoretical Inference, Scientific Instruments, and Testimony. *Theoretical Inference.* I may infer that there was a dog in my kitchen based on seeing muddy dog prints on a kitchen floor. *Scientific Instruments.* I may build a dog detector. Such a device may have a video camera and a computer on one end, and a green light on the other and be rigged up in such a way that the green light comes on when and only when there is a dog in front of the camera. I may see the green light, but not the dog, and nonetheless come to believe that there is a dog in front of the

camera of the detector. *Testimony.* A friend of mine may be stationed at the coffee house down the street sees a dog and calls me on the phone to tell me about it.

I have examined the debate between realists and idealists for the purposes of getting a bead on what the strength of dependence is that is at stake in arguing over mind-dependence. For realists and idealists to have anything to disagree about, the notion of dependence in question must be something stronger than mere material dependence. I have discussed two kinds of dependence stronger than material dependence. The first, which I called strong dependence, is counterfactual dependence. The second, which I called weak dependence is dispositional dependence. I redescribed these notions in terms of causal interactions. A property is strongly mind-dependent if there are causal interactions between it and minds, such that the property would not be instantiated unless there are minds. A property is weakly mind-dependent if the property, while being instantiable independently of minds, enters into casual interactions only with minds.

5.2 The side-ways view: Scientific Room for Idealism and Realism

> Hilary Putnam writes:
>
> I agree with Rorty that we have no access to "unconceptualized reality." As John McDowell likes to put it, you can't view your language "from sideways on" in the way that the idea of looking at one's language and the world and comparing the two suggests. (1994, p. 297)

McDowell is wrong. We can view the relation between mind and world from sideways on. We won't, of course, do it by stepping out of our minds to get a different view: we won't be able to see anything without thereby seeing it. But we can get a sideways on view of mind and world in much the same way an optometrist gets a sideways view of eye and chart: by viewing someone else's eye. Eyes are no less objects of empirical study for being among the things we use in our observations. Likewise for minds.

In this section I address the empirical and scientific warrant for positing both real and ideal kinds, that is, properties that are and properties that are not mind-dependent. This chapter is dedicated to shedding some light on the question of how you know when something is objective or subjective. The answer I offer is: you do the science. This will, of course require much unpacking. Among the many problems that such a project faces is the philosophical prejudice that such questions are necessarily beyond the grasp of science, that they are necessarily "metaphysical" in the pejorative sense of the term introduced and deplored by the logical positivists. Positivist resistance to realism, for instance, manifests itself as the belief that nothing can be added to an existence claim by asserting in addition that the entities in question are "real". My approach to such claims begins by seeing issues of metaphysics as continuous with issues in the natural sciences. All I mean by 'metaphysics' is the investigation of what things are. In this sense, then, scientists are in no less the metaphysical business than David Lewis or Jaegwon Kim.

My investigation of what is and isn't mind-dependent begins by recognizing mental states, thoughts and experiences as empirical objects, that is, objects of study by natural sciences on par with quarks and frogs. Thoughts are empirical objects, so are frogs. To say that something depends on something else in a metaphysically substantive way is to make an empirically testable claim about, say, explanatory or causal priority of one thing over an other. It is to say that there is a dependence above and beyond mere material equivalence. That the sea is salty is materially equivalent to Bill Clinton's being from Arkansas. But that is not a metaphysically substantive claim about the origins of the existence of the salty sea. In contrast, it is metaphysically substantive to say that the saltiness of the sea depends on the erosion of the adjacent landmass.

Recasting things in this way, even what is 'really real' is something that can meaningfully be talked about. And it can be shown to be distinct from being accidentally true, and it is something that we can have genuine confirmation of. Just as we can know whether frogs depend on nitrogen or the other way around, so we can know whether frogs depend on thoughts about frogs or the other way around. We have genuine access to what the representations are, and what things do and don't depend on being represented for their existence. Whether something is an objective or a subjective is something that we can have knowledge of: scientific knowledge.

My aim below is to present arguments that there are scientific grounds for positing that creatures represent mind-independent properties as well as mind-dependent

properties. Further, I argue that these are quite compatible situations, that is, that the reasons for thinking e.g., that creatures represent mind-dependent properties are not so strong that we must posit that those are the only properties that creatures are able to represent. Much of the burden of showing that mind-dependent properties are represented by organisms was taken up in chapters 2 and 3. Here I am especially concerned to show that these are not the only properties that organisms need to represent. I begin by considering arguments that a scientific view of the representational capacities of organisms leads to anti-realism. In this section I want to begin by discussing empirical considerations that are alleged to count against realism. The first is an old insight. It begins with noting the enormous amount of processing and mechanism that mediates between the distal stimulus and the eventuating of a percept. It is then portrayed as a small step from this insight to begin boggling at how we can come to perceive reality at all, it being so far away. Huw Price describes the position as

> the observation—naturalistic, note, though Kantian in spirit—that in virtue of the processing that goes on in perception, experience is necessarily removed from reality. While the transcendental perspective may be dismissed as ultimately incoherent—the product of an untenable metaphysical externalism—this naturalistic analogue is much more tenacious. In the end, indeed, it isn't "primitive metaphysics" but simple physiology which teaches us that Kant was right: what we get from our sensory apparatus depends on quite contingent features of our physical construction, as well as on the nature of the external world. Arguably, the same is true of our entire conceptual apparatus, but certainly it is true of experience. This product of a sideways-on scientific perspective is not a kind of comatose version of transcendentalism, but a plausible first-order theory about the way in which our brains are linked to their environment. Nor is it a kind of philosophical opening bid, which we can abandon on the grounds that it causes problems elsewhere in

> philosophy. To all intents and purposes it is a fact of modern life, within the constraints of which philosophy must operate. (forthcoming. Text cited from draft on author's web site.)

A similar kind of point is put forward by Rick Grush, describing a rabbit as a putative representer of its world:

> The rabbit, I know as theorist, has no immediate access to the world outside its skull. It is a brain in a small boney vat. But this brain constructs a world, and most of the time if not always, takes this internal construction for the world. (personal communication)

We should be cautious about claims about the rabbit having no immediate access to the world out side of its skull. What would it take for it (qua rabbit known to theorist) to have immediate access to the world outside? Would it count to take the brain out of the skull and, say, smearing it up against a pile of carrots? I worry that Grush (and perhaps Price) is moving from the fact that we access the world in virtue of representing it to a view that our representations some how get between us and the world, and thus prevent us from having immediate access to the world. But why shouldn't being representation producers and consumers be precisely the same thing as having immediate access to a world, at least part of which is mind-independent? Whatever access to a world amounts to, it should be something for which a distinction between mediate and immediate access at least makes sense. In order for the denial of immediate access to make sense, that which it denies must make sense. But the idea of the immediate has not yet been given a sense.

More can be said, of course, in favor of the kind of position that Grush alludes to—one in which empirical considerations concerning the relations of organisms to their environments leads to thinking of their represented worlds in idealist anti-realist terms. I turn now to examine further arguments that an empirical understanding of the relations between organisms and their environments leads naturally to viewing them as not representing a mind independent world, but instead, constructing a world of their own.

Von Uexkull describes an organism's represented environment as an *Umwelt*. The general point may be illustrated as von Uexkull does, by looking at the Umwelt of the tick. The tick's Umwelt consists of three receptor cues and three effector cues (Von Uexkull 1934, p. 12). A tick lingers in the leaves of trees until it detects butyric acid, a chemical emitted by mammals. Upon detecting some butyric acid the tick drops from its leafy perch, hopefully (for the tick) onto a mammal below. After the drop, the tick senses that contact has been made, which triggers a fury of scurrying. The third stimulus is a temperature cue that elicits burrowing into the flesh of the mammal for the bloody meal therein. As Von Uexkull puts the point, the effective environment, or 'umwelt', of the tick is constructed: "out of the vast world which surrounds the tick, three [stimuli] shine forth from the dark like beacons, and serve as guides to lead her unerringly to her goal"(op cit., p. 11). The tick thus operates on a pretty strict need-to-know basis. Many have taken this kind of point to generalize to all creatures great and small, thus giving the view that different animals perceive and cognize a "relevant-to-my-lifestyle world, as

opposed to a world-with-all-its-perceptual properties" (Churchland et al 1994, p. 56). Mandik and Clark (in review) call this view the 'thesis of selective representing'. The view includes the notion that, as Clark (1997) puts it, "aspects of real-world structure which biological brains represent will often be tightly geared to specific needs and sensori-motor capacities" (p. 173). The question I address in this section is whether the thesis of selective representing leads to the view that the way organisms represent the world is incapable of getting at the way that the world really is. Does it lead to the view, expressed by Varela et al, (1991) that the world as perceived and conceived by organisms, including human scientists, is in a non-trivial sense mind-dependent?

In a recent article, Anthony Chemero (1998) argues against realism. The crux of Chemero's argument moves from ideas about selective representing to the anti-realist conclusion that different ways of representing bring about the existence of different worlds that are represented.

Chemero's argumentative strategy may be unpacked further as follows. According to the thesis of selective representing, organisms' representational apparatuses operate on a pretty strict need-to-know basis. The way an organism represents the world is the result of a quick and dirty solution to a problem created by the special circumstances of the organisms' biological needs. Different niches give rise to different species-specific representational schemes. If it is safe to assume that both gibbons and goldfish represent the world, then it is also safe to assume that they represent the world in

radically different ways. A gibbon may represent the world as having good branches to swing from whereas whatever goldfish represent surely doesn't include swinging from tree branches.

Chemero makes the move from there being multiple species-specific representational schemes to there being multiple mind-dependent worlds brought about by these different representational schemes. As Chemero sees it, our human way of representing the world—including our scientific representational institutions like physics and all the rest—themselves are as much quick and dirty need-to-know solutions to a biological problem constrained by the biological peculiarities of our species. Our ways of representing are as different from gibbons' ways as the gibbons' ways are from the goldfishes' ways. Why, Chemero asks, privilege our way as the one way that gets it right? Why privilege our scientific ways of representing the world as the representations that represent the way the world really is?

There are two initial ways in which Chemero's argument goes wrong. The first way involves the supposition that the different ways of representing the world are in conflict—that they somehow constitute disagreement. Let us suppose that organism *X* represents only varying temperatures and that organism *Y* represents only varying concentrations of sulfur. The organisms are not disagreeing. It is not like *X* represents the presence of temperature and *Y* represents the absence of temperature. And since *X* and *Y*

are not in disagreement, it is entirely consistent to maintain that both *X* and *Y* represent the way the world really is.

Now there is a way to import disagreement into the situation of *X* and *Y*, but such a way is entirely illicit. One may attempt to redescribe the situation by saying that *X* represents the world as containing *only* temperature and saying that *Y* represents the world as containing *only* sulfur. Thus *X* and *Y* are representing the same thing in contradicting ways. The key point here is that such simple organisms simply are not equipped to represent the world as a whole and predicate of it the presence of only varying degrees of sulfur. They simply do not have the conceptual resources to pull this off. Instead the situation is akin to one person saying that dogs are furry and another person saying that dogs have four legs. While each person is saying different things, the situation need not be one of disagreement, thus it is entirely consistent to maintain that both people represent the way dogs really are.

The illicit importation of disagreement arises by treating as interchangeable "*X* only represents *Y* as *Z*" and "*X* represents *Y* as only *Z*." One must guard against such a maneuver. One must not treat as interchangeable the phrases like "The only things that George thinks about are cheese burgers" and "George thinks that only cheese burgers exist". In the first situation, George need employ only the concept of cheeseburgers. But in the second situation, George needs in addition to the concept of cheeseburgers, the

concepts of existence and negation. Chemero does not adequately guard against such an illicit move. Chemero writes:

> Because the needs of one type of animal can be are so [*sic*] different from those of another, the perceptual systems that result will constitute the world in very different ways, as full of barbecues and highways and myriad other things for humans, but, for example, as containing only three things—what we see as butyric acid, pressure and temperature changes—for ticks (see von Uexkull, 1934, p. 10). (paragraph 15)

In the above passage, the assertion that ticks represent the world as containing only butyric acid, pressure, and temperature changes is unwarranted and may not be inferred from the mere fact that the only aspects of the world ticks are capable of representing are butyric acid, pressure, and temperature changes. It is one thing to say that ticks represent only *X*, *Y,* and *Z*. It is an entirely different thing to say that ticks represent the world as having only *X*, *Y,* and *Z*. The latter case is what is needed for the tick's representations to be in conflict with ours. But the former case is all that the thesis of selective representing is committed to, and the former case is consistent with realism.

The second way that Chemero goes wrong is by a fallacious supposition of the exclusivity of functions. The fallacy is to infer from the premise that the function of organisms' representational schemes is to get by, to the conclusion that the function of organisms' representational schemes is not to represent the way the world really objectively is. Chemero writes:

> [G]iven the way evolution works, we should not think of the perceptual systems (or any parts of animals) as ideal solutions to problems posed by the environment. Instead, animals that survive and reproduce are

> those that do well enough to find food and so on. So, there is no reason to assume that any particular animal's perceptual system gets the world, as it is independently of thought, just exactly right; they all do only well enough. (paragraph 15)

Continuing on this theme, and drawing on Clark 1997, Chemero writes:

> Consider that Clark argues that "higher thought," the kind exhibited in mathematical and scientific theorizing, depends on the scaffolding provided by public language. He also suggests (pp. 211-13) that language is adapted to the way our brains worked pre-linguistically; human language, that is, is adapted to and built upon action-oriented representations. But, as we have seen, these representations are biased by pressures to fulfill human needs throughout evolutionary history. And if the foundation on which language is built is biased, it is overwhelmingly likely that language itself is similarly biased. So if physics and other sciences depend upon our language-using abilities (and Clark argues that they do), they have no claim on being reflections of the world-in-itself. (paragraph 19)

Chemero's passages echo suspicions that have been around for a while. For example Patricia Churchland (1987) urges "Looked at from an evolutionary point of view, the principle function of nervous systems is to get the body parts where they should be in order that the organism may survive. . . . Truth, whatever that is, definitely takes the hindmost" (pp. 548-549).

I offer in response that two different functions can be compatible: an organism can be tightly fit into a particular and peculiar niche and represent the way things really are. To suppose otherwise, that is, to suppose the exclusivity of the two functions is like arguing that the function of a stop sign is not to get cars to stop because the function of a stop sign is to help prevent car accidents (See also Grush and Mandik (in press)). A tick

may represent just what it needs to get by: concentrations of butyric acid etc. But this is entirely compatible with representing the way things really are: as being concentrations of butyric acid, etc.

The rhetorical device Chemero employs and that is worth pointing out is the way that he moves between the phrases "The way X represents the world" and "X's world". This rhetorical device paves a smooth passage for anti-realism for it makes it seem that there is a world for each way of representing—that different representational schemes are different ways of world-making. Admittedly, such language is encouraged by the introduction of appealing to Umwelts in the first place. Nonetheless, one must avoid the view that makes idealism rest on the following tautology: the only world that we represent is a world that is represented by us. Now, of course, the world represented by us is representation dependent in at least this sense: it depends on being represented by us for its being represented by us. But, as discussed in section 5.1, this can't be what the realist and anti-realist are disagreeing about. I mention this point not attribute it to Chemero, only to point out the dangerous proximity between the tautologous version of anti-realism and the rhetoric employed in these discussions.

Chemero's attempt to derive non-realist conclusions directly from the notion of selective representing fails. The thesis of selective representing does not work against a realism about the entities posited by our scientific explanations. Instead, the selective representing story depends on taking seriously the mind-independent existence of what is

posited. Part of the scientific understanding of the tick's Umwelt includes the notion that the sorts of things that the tick is responsive to—temperature and butyric acid—are the sorts of things that exist independently of tick's responsivity. The scientific story licenses saying that the tick is responsive to stuff that would be there even if the tick were not. Chemero wonders what reason the scientific realist has for supposing that human Umwelten capture the way the world really is. The reason the realist provides is the same reason for supposing that tick Umwelten capture the way the world really is: the world really does contain concentrations of butyric acid and ticks are responsive to those concentrations.

In advocating this kind of picture, a picture in which not just human scientists, but lowly ticks are capable of representing the way the world really is, I do not intend the case to exhaust the scientific options for describing the Umwelten of various creatures. While it may sometimes be the case that a creature's mind is in touch with aspects of a world independent of its mind, there are also instances in which the creature's Umwelt includes mind-dependent properties, as argued in chapter 2.

Having sketched how we may find ourselves attributing an objective content to the state of an organism, I want to turn now to considerations favoring instances in which it may be warranted to attribute subjective contents. One way we may do this is by keeping our attention fixed on the examples of thermometers and the detection of temperatures. How might one find oneself deciding that what seemed initially to be in

the business of detecting the objective property of temperature was actually in the business of detecting some creature-dependent property? Answers to these sorts of questions were discussed at some length in chapter 2. In addition to the sorts of examples given there, another class of examples of subjective contents comes from the literature on Gibsonian/ecological approaches to psychological understanding. The gist of these examples will involve psychological states that have as their contents things that depend on the organism itself. For instance, if I see something as too large for me to eat, there is a real sense in which the property I attribute to that object depends on me. Whether something instantiates the property of being too large for me to eat depends on facts about me—my size and metabolism etc.

Gibson 1979 gives several examples in which the organism itself, or parts of the organism, or properties that depend on the organism figure in the contents of visual experience. (Compare the discussion of Tanky from chapter 2.) In a discussion of these points, Bermudez 1998 calls such contents 'self-specifying information'.[8] The first example in which self-specifying information shows up in the contents of visual experience concerns the boundaries of the visual field. In the case of humans, the visual field contains less than half of the viewer's visible environment. The boundary of the visual field constitutes a boundary between the seen and unseen world that is movable at will and supplies information to the viewer about her location in and trajectory through

[8] The following draws heavily from Bermudez's (1998) discussion.

that world. Parts of the viewer's body (especially the nose and cheeks) define the boundary of the visual field, and some parts, the nose in particular, is itself seen in the visual field. The occlusion of other seen objects by the nose supplies information about degrees of proximity of objects to the viewer. Another important proximity clue involves the amount of the visual field an object occupies: as objects move away from the viewer, the occupy less of the visual field. This fact allows for an interesting way of visually distinguishing between parts of the subject's body and the rest of the seen world. The viewer's hand or foot can only move so far from the viewer, being typically and hopefully attached to the viewer, and thus there is a limit on how little of the visual field it can be restricted to. In contrast, other objects, being unattached, may move indefinitely far away from the viewer and thus occupy an indefinitely small portion of the visual field. This difference may also facilitate a clue as to what objects, though not a part of the organism's body, are at least attached to it or grasped by it. A stone held in my hand cannot occupy an indefinitely small portion of the visual field. Only if it is let go may it shrink into the distance. Another source of self-specifying information is optical flow. As a subject moves forward, the shapes in the visual field flow from the center to the boundaries. The center from which the flow emanates specifies the direction in which the subject is moving. These examples drawn from Gibson constitute rich sources of information available to the perceiving organisms. The step from the availability of this information to subjective representation is a small one. If we suppose that organisms have states that function to carry that information, and that information includes

information about the relation of the environment to the subject, then we have on our hands representations that are perspectival/egocentric/subjective.

5.3 The private and the public

Discussion in this chapter so far has centered on the question of how one might, from a naturalistic perspective, go about answering certain "how do you know?" questions concerning objectivity and subjectivity. In particular, the questions of focus have concerned how in a scientific context one might assign representational contents to the states of organisms in accordance with either realist or idealist conceptions of the relation of mind and world. I want to turn now to a closely related topic at the interface of objectivity and epistemology, the topic of the relation of the objective/subjective distinction to the epistemological distinction between publicly accessible and privately accessible facts. This discussion will draw upon ideas set forth earlier in the chapter, and will serve, among other things, to highlight some differences between the correspondence theory of epistemic objectivity and its competitor, the consensus theory.

Frequently associated concepts in the philosophical literature are the associations of subjectivity with privacy and objectivity with publicity. I have heretofore sketched accounts of objectivity and subjectivity but not explicitly addressed privacy and publicity. Privacy and publicity are oft cashed out in epistemological terms, thus it is most fitting that I have postponed their discussion until the chapter on the epistemology of objectivity. In brief, the epistemological construal of these notions is as follows. Something is private if epistemic access to it (knowledge of it or evidence about it) can only be had by a single person. Thus on certain Cartesian conceptions of mentality,

thoughts and experiences are private: only I know what, if anything, I am thinking and feeling. Thus the problem of other minds becomes especially acute: how can I know that other individuals are hosts of any mentality whatsoever and not merely mindless automatons? Attributing privacy to sensory qualities leads to the problems surrounding the knowledge argument discussed in chapter 3.

In contrast, something is public if epistemic access to it (knowledge of it or evidence about it) can be had by multiple individuals. Thus the Eiffel tower is a public entity and the weight of a cubic centimeter of water is a public property. A notion closely related to that of publicity is the notion of consensus or agreement. What makes a quality like mass public is that we can employ methods whereby we reach agreement on the mass of some individual. We can each weigh a sample of gold and agree that it has mass of *n* grams. Suppose, in contrast, that I assert that the experience of drinking orange juice right after brushing my teeth with mint toothpaste is sublime. What method would you employ to verify for yourself the sublimity of that experience? You may drink some of your own orange juice after brushing your own teeth and find that your experience is far short of capturing the sublime. But that's your experience, not mine. How do you come to know that my experience is sublime? You take my word for it, and trust that I wasn't just trying to trick you into wasting some orange juice. But, the story goes, while you have to settle on taking my word for it, only I can really know. The methods of rational agreement employed in the determination of the masses of objects have no place in

determining the sublimity of experiences. Thus are the notions of publicity and consensus alleged to go hand in hand.

For some philosophers, the association of objectivity and consensus is so strong that they see an identity: the objective truths just are the ones we agree on. The correspondence theorist rejects this identification. But for the correspondence theorist an obligation stands to at least explain what the private/public distinction has to do with the subjective/objective distinction. Before taking on the topics of the privacy of subjectivity and the publicity of objectivity, it will be helpful to briefly recap the basic view so far.

If there is a property *P* that we can represent with some representation *R*, then there must be at least one kind of causal interaction that *P* enters into, namely one that relates *P* to *R* without which, *R* could not represent *P*. If the only kinds of interaction that *P* enters into are interactions with representations, then *P* is metaphysically subjective, at least in the weak sense of being a lovely property (though the interactions *P* enters into may be such to make it strongly mind-dependent). If, however, *P* enters into interactions with another property *P'* and *P'* is not a representation, then *P* is metaphysically objective (it is neither strongly nor weakly mind-dependent). These ideas connect up with privacy and publicity as follows.

If you and I agree that *a* is *F* then there must be at least two representations that enter into interactions with *F*: your representation of *F*, *R*1, and my

representation of *F*, *R*2. Suppose that the only laws that *F* entered into were those that related it to the representations *R*1 and *R*2. *F* would be subjective. But there would be agreement on *F*. Thus agreement is not sufficient for objectivity and objectivity is not necessary for agreement. Suppose, however, that there is a property that is maximally subjective, that is, it enters into only one causal interaction and that interaction relates it to one representation. If a property is maximally subjective, then it will be private.

I turn now to flesh these ideas out further. Common scientific practice has it that lines of converging evidence are indicative of the truth of a hypothesis. The more lines of evidence that converge, the more likely the hypothesis. This point holds not only at the level of theory, but at the level of data. How does one go about deciding that a putative datum is not just an artifact? For example, if using one kind of microscopy, one observes a bulge in a cell wall. One may wonder whether that bulge is just an artifact of the particular preparation. One tries multiple preparations to see if the bulge appears in more than one of them. (For an extended discussion of precisely this sort of issue as it arises in cell microscopy, see Bechtel 1994.)

Why is it that multiple lines of converging evidence are taken as indicative of objectivity? And how can this be explained in terms of the correspondence theory, that is, without following the consensus theorist in making consensus constitutive of objectivity? To see the how these ideas can be worked out, it will be useful to return to

the psychosemantic project of determining the representational contents of states of an organism.

Let us presume that with respect to frogs, flies are mind-independent: with respect to our science, at least, flies are frog independent, so let's just grant that. Let us also presuppose the old story about the fly detector (Lettvin *et. al.* 1959): that there is some event type that occurs in the frog's brain that is relatively well correlated with the presence of flies. Call that brain event *B*. Following Godfrey-Smith (1998, p.250) let us call the conditional probability that a fly is present given *B* 'reliability' and the conditional probability that *B* will occur given the presence of a fly 'sensitivity'. We can define the frog's hunger as the amount of time that has elapsed since it last ate a fly. Observing the frog and its brain activity, we note the following. The frog's sensitivity to flies increases as a function of its hunger. The conditional probability that the presence of a fly will activate *B* is much higher when the frog is hungry.

Given these presumed facts let us question whether *B* really functions to detect/represent flies (instantiations of the property of being a fly) or some other and perhaps mind-dependent (frog's-mind dependent) property.

Ex hypothesi, the property of being a fly is a frog-independent property. But consider the property defined as being a fly-when-a-frog-is-hungry (FWAFIH). FWAFIH is pretty obviously frog-dependent. And while the frog is relatively sensitive to

flies in virtue of brain activity *B*, the frog is even more sensitive to FWAFIH's. And arguably, what matters to frog survival is not so much the detection of flies simpliciter, but instead detecting flies when its hungry. Is it so large a leap to think that the frog is in the business of detecting FWAFIH's?

Maybe it is. Having the function of detecting flies when a frog is hungry is distinct from the function of detecting the presence of fly-when-frog-is-hungry. But how to draw the distinction? How would one go about deciding over fly or fly-when-frog-is-hungry? A handle on this question can be gotten by deciding an analogous question asked regarding scientific instruments instead of sensory systems. Let us return to the example of thermometry discussed in chapter 2. How would one decide that an instrument, say a thermometer, were detecting an objective property temperature instead of the thermometer-dependent property *temperature-of-thing-thermometer-is-in/near*?

The key is noting which properties are explanatory, and which count among what needs to be explained. If what needs to be explained is why a thermometer has a certain reading when it is in a substance, then the elements of an explanans that are restatements of what was said in the explanandum are strictly superfluous. Circularity is the death of explanation. "Grass is green because grass is green" is so obviously unexplanatory because it is so obvious that the properties picked out in the explanans are all and only the properties picked out in the explanandum. Molier's famous and funny example of explaining the fact that opium makes people sleepy in terms of positing a dormative

virtue of the opium is both famous and funny because citing a dormative virtue of opium only restates what was known already: that opium makes people sleepy. What needs to be explained is the why the thermometer goes up when its placed in a substance. Invoking the property of temperature is a candidate explanation. But consider the property temperature-of-substance-thermometer-is-in. Note the overlap between the explanans and the explanandum. We already know that the thermometer is in the substance. Invoking the thermometer dependent property bears no explanatory load over and above that born by the invocation of the thermometer independent property. We are interested in the property that explains why, when the thermometer is in the liquid, the mercury goes up. We abstract away from the instrument. When you have a fever, you instantiate a property that is detected by the thermometer, and, let us grant, a property that depends on thermometers, but these are not the same because of the differential explanatory statuses of the two properties.

These remarks allow us to see why consensus has more bite in understanding objectivity on the correspondence view. The agreement of instruments only makes sense on the supposition that they are in the business of detecting properties independent of the instruments. Likewise, using one instrument to calibrate another (and how else can instruments be calibrated?) only makes sense on the supposition that the two different instruments are tapping into something independent of the two instruments. If what each instrument was in the business of detecting was a property dependent on that instrument, then each separate instrument would be detecting its own separate property and there

would be no common measurement that they could be said to be agreeing on (and there would be no room for the notion of calibration to get a grip).

This argument is distinct from the common one that treats the objectivity of what is measured as an inference to an explanation of the agreement of the instruments. The common argument assumes agreement, and then posits the objectivity of what is measured as an inference to a best explanation of why there is agreement. What I am proposing is different from this common argument. What I am proposing is that the very notion of agreement itself presupposes objectivity (in the sense of instrument-independence of the properties measured). If what is measured is instrument dependent, then the notion of agreement cannot grab hold because where there are different instruments, then there are different instrument-dependent properties that each is tracking.

Thus we have a sketch on how one might go about deciding in the favor of the 'realist' option for attributing contents to the states of the frog. What I mean by 'realist' in this context is the option of deciding that the property some detector (or instrument) detects or measures is a property that is instantiated independently of the detector. The sketch shows how notions of objectivity hook up with notions of consensus or agreement without being identical to those notions. Indeed, part of the notion of agreement between, say, measuring instruments presupposes a metaphysically objective thing for them to be in agreement about.

Just as the psychosemantic explication of objectivity shows up the connections to publicity, we may show the relations of subjectivity with privacy. A fact or property is private insofar as epistemic access to it is restricted. This notion is cashed out in terms of the property's entering to causal interactions only with representations of it. Such a psychosemantic reconstruction of the notion of privacy is illustrated nicely by the following example. Dennett (1997, p. 637) describes an old spy trick of tearing a piece of paper, or in the case of the Rosenbergs, a Jell-O box, to create an un-forgeable identifier. The matching halves are so informationally complex, that the only way of identifying one, is by identifying the other. Dennett suggests a similar match between some sensory states and their corresponding worldly properties, and further suggests that this feature gives rise to a certain amount of our intuitions regarding the privacy and ineffability of subjective experience.

There is a very real sense in which the one piece of the Jell-O box depends on another, the one piece would not have turned out this way if the other hadn't turned out that way. There is a counterfactual dependence of the one on the other. There is also an analogue of the weaker sense of mind dependence of the one piece on the other: the property of being that piece brings with it no other properties except those had by the other piece. We can treat one piece as a detector of the other. There is a nice analogy between the Jell-O-box pieces thought of as involved in the process of detection and the notion of an experimental artifact: in both instances, what is detected is in part a product

of the process by which it is detected. We can also imagine that there are representational states that match their representeds in the way that the Jell-O box pieces match each other. The properties detected by such a sensory system could not be detected by anyone else lacking such a system and would thus be effectively private for the owner of the sensory system.

Given the context of teleosemantics, the question arises of how a uniquely instantiated property can be represented. The puzzle arises if we think of teleosemantics as requiring multiple instances of a property figuring in the environment of a population subjected to the forces of natural selection. How could a teleosemantics account for the representational contents of uniquely instantiated representational systems? If we see the representations of part of a scheme, such as a slider along a scale representing temperatures, then we can imagine how that scheme can come to have the function of representing a continuous range of temperatures even though there are particular temperatures within that range that have never been represented. More generally, the idea is that what would be selected for would be a relatively general purpose scheme of representation, say the representation of visual patterns, that comes to, in virtue of the idiosyncrasies of a particular subject, represent specific patterns that had never been represented in the organisms' ancestral history.

Bibliography

Akins, K. (1996). Of sensory systems and the 'aboutness' of mental states. *The Journal Of Philosophy* 93: 337-372.

Alter, T. (1998). A limited defense of the knowledge argument. *Philosophical Studies* 90, 35-56.

Anscombe, G. (1957). *Intention.* Oxford: Basil Blackwell.

Audi, R. (1992). Objective/subjective. *A Companion to Epistemology*. J. Dancy and E. Sosa. Eds. Oxford, Blackwell: 309-310.

Bach-y-Rita, P. (1972). *Brain Mechanisms in Sensory Substitution.* New York and London: Academic Press.

Bechtell, W. (1994). Deciding on the data: Epistemological problems surrounding instruments and research techniques in cell biology. *PSA 1994*. Volume 2. pp. 167-178

Bell, D. (1992). Objectivity. *A Companion to Epistemology*. J. Dancy and E. Sosa. Eds. Oxford, Blackwell: 310-312.

Benacerraf, P. (1965). What numbers could not be. *Philosophical Review. 74*: 47-73.

Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy. 19*: 661-79.

Berkeley, G. (1710). *A Treatise Concerning the Principles of Human Knowledge: Wherein the chief causes of error & difficulty in the sciences, with the grounds of skepticism, atheism, and irreligion, are inquired into.*

Berkeley, G. (1713). *Three Dialogues Between Hylas and Philonous, In Opposition to Skeptics & Atheists.*

Bermudez, J. (1998). *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.

Brewer, B. (1996). Thoughts about objects, places and times. *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind* (Proceedings of the British Academy, Vol. 83), C. Peacocke (ed.). Oxford: Oxford University Press.

Bülthoff, H., and Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Science, 89,* 60-64.

Campbell, J. (1996). Objects and objectivit. *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind* (Proceedings of the British Academy, Vol. 83), C. Peacocke (ed.) . Oxford: Oxford University Press.

Chemero, A.(1998). A stroll through the worlds of animats and persons: A review of Andy Clark's *Being There*", *Psyche*, 4(14). http://psyche.cs.monash.edu.au/v4/psyche-4-14-chemero.html

Churchland, P.S., Ramachandran, V., et al. (1994). A critique of pure vision. In C. Koch and J. Davis (eds.) *Large-Scale Neuronal Theories of the Brain*. Cambridge, MA: MIT Press.

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Dennett, D. (1991). Lovely and suspect qualities, (commentary on Rosenthal, " The Independence of Consciousness and Sensory Quality"), in Villanueva, E. ed., *Consciousness,* Atascadero, CA: Ridgeview. Cited text from electronic version available on Dennett's web site.

Dennett, D. (1998). Quining qualia. *The Nature of Consciousness*. Ned Block, Owen Flanagan, and Guven Guzeldere, eds. Cambridge, MA: MIT Press.

Deutsch, M. (unpublished). Subjective physical facts.

Dretske, F. (1988). *Explaining Behavior*. Cambridge, MA: MIT Press.

Dretske, F. (1995). *Naturalizing the Mind.* Cambridge, MA: MIT Press.

Dummett, M. (1975). The philosophical basis of intuitionist logic". *Studies in Logic and the Foundations of Mathematics* 80: 5-40.

Eichenbaum, H., Stewart, C., & Morris, R. (1990). Hippocampal representation in spatial learning. *Journal of Neuroscience, 10*, 331-339.

Evans, G. (1982) *The Varieties of Reference*. John McDowell (ed.). New York : Oxford University Press.

Evans, G. (1985a) Molyneux's Question. In *Collected Papers.* New York : Oxford University Press.

Evans, G. (1985b) Things without the mind. In *Collected Papers.* New York: Oxford University Press.

Feigenbaum, J. and Rolls, E. (1991). Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology*. *19 (1),* 21-40.

Field, H. (1989) *Realism, Mathematics, and Modality*. Oxford: Blackwell.

Fodor, J. (1998). *Concepts*. Oxford: Oxford University Press.

Foss, J. E. (1993). Subjectivity, objectivity, and Nagel on consciousness. *Dialogue* 32: 725-36.

Gallistel, C. (1980). *The Organization of Action: A New Synthesis*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gauker, C. (1995). *Thinking Out Loud*. Princeton, NJ: Princeton University Press.

Gibson, J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Godfrey-Smith, P. (1998). *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.

Goodman, N. (1976). *The Languages of Art*. Indianapolis: Hackett.

Goodman, N. (1978). *Ways of Worldmaking*. Indianapolis: Hackett.

Grush, R. and Mandik, P. (in press). Representational parts. *Houston Studies in Cognitive Science*.

Grush, Rick (1997). The architecture of representation. *Philosophical Psychology* 10(1):5-23.

Held, R. and Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology* 56(5):872-876.

Hensel, H. (1982). *Thermal Sensations and Thermoreceptors in Man*. Springfield, IL: Charles Thomas.

Kirsh, D. (1996). Today the earwig, tomorrow man? *The Philosophy of Artificial Life*. Margaret Boden, ed. Oxford: Oxford University Press.

Kosslyn, S.M. (1994). *Image and Brain*. Cambridge, MA: MIT Press.

Kripke, S. (1972). *Naming and Necessity*. Cambridge, MA: MIT Press.

Lettvin, J., Maturana, H., McCulloch, W., and Pitts, W. (1959). What the frog's eye tells the frog's brain." *Proceedings of the IRF*.

Lycan, W. (1996). *Consciousness and Experience*. Cambridge, MA, MIT Press.

Mandik and Clark (in review). Selective representing and world-making.

McGinn, C. (1991). *The Problem of Consciousness*. Oxford: Basil Blackwell.

McGinn, C. (1995) Consciousness and space. *Journal of Consciousness Studies.* 2: 220-230.

Milner, A. and Goodale, M. (1995) *The Visual Brain in Action.* Oxford: Oxford University Press.

Moore, A. (1997). *Points of View*. Oxford: Clarendon Press.

Nagel, T. (1974) What is it like to be a bat? *Philosophical Review 83*, 435-50.

Nagel, T. (1986) *The View from Nowhere.* Oxford: Oxford University Press.

Newell, R. (1986). *Objectivity, Empiricism, and Truth.* New York, Routledge & K. Paul.

O'Keefe, J. (1996). Cognitive maps, time and causality. *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind* (Proceedings of the British Academy, Vol. 83), C. Peacocke (ed.). Oxford: Oxford University Press.

O'Keefe, J. and L. Nadel (1978) *The Hippocampus as a Cognitive Map.* Oxford: Oxford University Press.

Pappas, G. (1995). George Berkeley. *The Cambridge Dictionary of Philosophy*. R. Audi. Ed. Cambridge: Cambridge University Press. 72-74

Pitcher, G. (1970). *A Theory of Perception*. Princeton, NJ: Princeton University Press.

Price, H. (forthcoming). Review of John McDowell's *Mind and World. Philosophical Books*. Text cited from draft on author's web site.

Putnam, H. (1994). *Words and Life.* Cambridge, MA: Harvard University Press.

Quine, W. (1960). *Word and Object.* Cambridge, MA: MIT Press.

Quine, W. (1980). On what there is. *From a Logical Point of View. Second Edition, revised* . Cambridge, MA: Harvard University Press.

Rawls, J. (1993). *Political Liberalism.* New York: Columbia University Press.

Rescher, N. (1997). *Objectivity : The Obligations of Impersonal Reason*. Notre Dame, IN: University of Notre Dame Press.

Rey, G. (1997). *Contemporary Philosophy of Mind.* Oxford: Blackwell.

Rorty, R. (1979) *Philosophy and the Mirror of Nature*. Princeton, NJ: Princeton University Press.

Schmitt, F., F. (1995). *Truth: A Primer*. Boulder, CO: Westview Press.

Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press

Shepard, R. N., & Cooper, L. A. (1982). *Mental Images and Their Transformations*. Cambridge, MA: MIT Press.

Stein, J. (1992). The representation of egocentric space in the posterior parietal cortex. *Behavioral & Brain Sciences, Vol 15(4)*, 691-700.

Strawson, P. (1980). Reply to Evans. In *Philosophical Subjects: Essays Presented to P. F. Strawson*, Z. van Straaten (ed.): 273-82. Oxford: Clarendon Press.

Strawson, P. (1959). *Individuals.* London: Methuen.

Taylor, C. (1979). The validity of transcendental arguments. *Proceedings of the Aristotelian Society. 79*: 151-65

Tye, M. (1995) *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition, 32*, 193-254.

van Fraassen, B., (1980). *The Scientific Image*. Oxford: Clarendon Press.

Varela, F., Thompson, E., et al. (1991). *The Embodied Mind*. Cambridge, MA: MIT Press.

Von Uexküll, J. (1934). A stroll through the worlds of animals and men. In *Instinctive Behavior,* ed. K. Lashley. New York: International Press.

Wrenn, C. (1998) Explanations of objectivity in mathematics. Paper presented at the 1998 meeting of the Southern Society of Philosophy and Psychology.