

Mental representation and the subjectivity of consciousness^{*}

PETE MANDIK

ABSTRACT *Many have urged that the biggest obstacles to a physicalistic understanding of consciousness are the problems raised in connection with the subjectivity of consciousness. These problems are most acutely expressed in consideration of the knowledge argument against physicalism. I develop a novel account of the subjectivity of consciousness by explicating the ways in which mental representations may be perspectival. Crucial features of my account involve analogies between the representations involved in sensory experience and the ways in which pictorial representations exhibit perspectives or points of view. I argue that the resultant account of subjectivity provides a basis for the strongest response physicalists can give to the knowledge argument.*

Introduction

Consciousness has been a real pain for philosophers and cognitive scientists interested in developing materialist theories of the mind. Particularly vexing is the *subjectivity* of consciousness: the subjective character of conscious experience. What subjectivity is and why it poses problems for physicalists emerges most vividly in discussions focusing on the work of philosophers Thomas Nagel and Frank Jackson. Nagel (1974) famously asked “what is it like to be a bat?” and urged that humans, not being bats, could never know. Most importantly in his criticism of the limits of physicalism, Nagel argued that no amount of knowledge of the objective physical facts about bat brains could yield knowledge of the subjective character of bat experience. What is known from the objective point of view cannot bridge the gulf to yield what is known from the subjective point of view: knowledge of *what it is like*. Drawing on similar intuitions, Jackson marshaled these themes in an attempt to falsify physicalism by way of his now famous *knowledge argument*. 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

^{*}This paper was the winner of the 2000 William James Prize of the Society for Philosophy and Psychology, awarded for the best paper by a graduate student presented at the Society’s annual meeting.

Pete Mandik, *Department of Philosophy, William Paterson University of New Jersey, 265 Atrium Building, 300 Pompton Road, Wayne, NJ 07470, USA, email: mandikp@wpunj.edu*

in learning a new fact upon having a red experience, Mary learns a non-physical fact. Thus the subjective character of experience is non-physical (Jackson, 1982).

There have been many physicalist responses to Nagel and Jackson over the years. Most are flawed and only a few explicitly address the topic of subjectivity. In this paper I develop a theory of subjectivity and use it in a defense of physicalism against the knowledge argument. Crucial aspects of my account involve explications of the notions of *perspective* and *point of view*. I cultivate an account of perspective as it arises in pictorial representations and apply the resultant account of pictorial perspective to mental representations. I contrast my account against accounts that analyze perspectives and points of view in terms of indexicals. Key aspects in my arguments for the superiority of the pictorial theory of subjectivity over indexical theories hinge on how the respective theories are able to deal with the issues raised in connection with the knowledge argument.

This paper is organized as follows. I begin with further discussion of Nagel and Jackson's original arguments with special attention paid to the role played by the notion of subjectivity. I then briefly discuss the main kinds of physicalist responses to Jackson's knowledge argument. After canvassing the main kinds of physicalist responses and their flaws, I launch into the articulation of my pictorial theory of subjectivity. I discuss crucial similarities and differences between my pictorial theory and indexical theories, and address several possible objections to the pictorial theory.

1. The subjectivity of consciousness and the knowledge argument against physicalism

Nagel's reflections on objectivity and subjectivity have been developed in several places, but here I focus on his 1974 article "What is it like to be a bat?" Nagel presumes that bats are conscious, that there is something it is like to be a bat. Things that lack conscious, like rocks, are things for which there is nothing it is like to be. Not only is Nagel confident that there is something it is like to be a bat, he is confident that it is radically unlike what it is like to be a human. This is plausible. Bats have sensory systems quite unlike ours; most notably they employ sonar in their perceptual interaction with their environment. Given the divergence between humans and bats, humans are incapable of knowing what it is like to be a bat. Imagination and extrapolation are of little help. Even imagining flapping around and hearing our own voices reflected off of the surfaces of objects serves only to tell us what it would be like to be *us* under altered circumstances—such imaginings bring us no closer to knowing what it would be like for *bats*. What we can know about bats are the deliverances of objective physical sciences: facts and properties knowable from multiple points of view—objective facts. The facts about bats beyond our purview are the subjective facts: facts known, if at all, only from the point of view of bats.

A crucial point to note, and one that I will return to several times in this paper, is that Nagel intends these remarks about perspectives and points of view to be about *types* and not *tokens*. Nagel writes:

I am not advertent here to the alleged privacy of experience to its possessor. The point of view in question is not one accessible only to a single

individual. Rather it is a *type*. It is often possible to take up a point of view other than one's own, so the comprehension of such facts is not limited to one's own case. There is a sense in which phenomenological facts are perfectly objective: one person can know or say of another what the quality of the other's experience is. They are subjective, however, in the sense that even this objective ascription of experience is possible only for someone sufficiently similar to the object of ascription to be able to adopt his point of view—to understand the ascription in the first person as well as in the third, so to speak. The more different from oneself the other experiencer is, the less success one can expect with this enterprise. (1974, pp. 441–442)

Nagel thus urges that what we are cut off from are facts about certain properties—universals or repeatables that are instantiable by multiple individuals. Nagel's arguments are not targeting the epistemic problem of knowing the minds of those that are merely numerically distinct from you, a problem that could arise between you and an otherwise identical twin. Nagel is instead concerned to point out a problem that arises in trying to comprehend a quite qualitatively distinct *kind* of mind.

Nagel casts his criticism as pointing out the limits of current physicalistic understanding, but he does not see it as falsifying physicalism. While he has no doubt *that* physicalism about subjectivity is true, Nagel says that we cannot currently understand *how* it is true. Jackson, in contrast, cultivates Nagelian themes into an argument against physicalism: the knowledge argument. Jackson's knowledge argument involves two thought experiments, one more famous than the other. The famous thought experiment features Mary, the physically omniscient scientist who only comes to learn what it is like to see red late in life. Less famous is the thought experiment that features Fred. Where Mary spends most of her life restricted to the black and white room and thus seeing far fewer colors than the average human, Fred is gifted with an unusually sensitive visual system and thus sees more colors than the average human. Where average humans might group a set of ripe tomatoes together as all being the same shade of red, Fred sorts the tomatoes into two groups that he calls red1 and red2. Fred assures us that red1 and red2 are as different from each other as yellow and blue. He sorts the same objects into the same groups trial after trial, giving evidence that the perceptual categories that he is tapping into are quite salient and robust. Fred's reports and behavior make us confident that what it is like to be Fred is different from what it is like to be a normal human. However, we do not know what it is like to be Fred. We may acquire all of the physical information about how Fred's nervous system works and what physical features of objects he is sensitive to, but this is insufficient to give us knowledge of what it is like. According to Jackson, however, there are other conditions that will suffice to give us knowledge of what it is like to be Fred. Jackson invites us to imagine that humans have their visual systems surgically altered to resemble Fred's, or upon Fred's death someone else receives a transplant of Fred's visual system. After these surgeries, people would come to learn something new—what it is like to be Fred. Prior to the surgery, however, the regular humans knew all of the physical facts about Fred. So in coming

to learn something new, they come to learn something non-physical. Thus physicalism is false since not all of the facts are physical facts.

Like Nagel, Jackson urges that these remarks about the subjective character of experience target a point about types of minds, not tokens. Jackson writes:

When I complained that all the physical knowledge about Fred was not enough to tell us what his special color experience was like, I was not complaining that we weren't finding out what it is like to *be* Fred. I was complaining that there is something *about* his experience, a property of it, of which we were left ignorant. And if and when we come to know what this property is we still will not know what it is like to *be* Fred, but we will know more *about* him. No amount of knowledge about Fred, be it physical or not, amounts to knowledge "from the inside" considering Fred. We are not Fred. There is thus a whole set of items of knowledge expressed by forms of words like "that is *I myself* who is ..." which Fred has and we simply cannot have because we are not him.

When Fred sees the color he alone can see, one thing he knows is the way his experience of it differs from his experience of seeing red and so on; *another* is that he himself is seeing it. Physicalist and qualia freaks alike should acknowledge that no amount of information of whatever kind that *others* have *about* Fred amounts to knowledge of the second. My complaint though concerned the first and was that the special quality of his experience is certainly a fact about it and one which Physicalism leaves out because no amount of physical information told us what it is. (1982, p. 132)

In other words, whatever indexical knowledge Fred has uniquely in virtue of being numerically distinct from everyone else is not the knowledge that Jackson sees as problematic for physicalism. Instead, the problem arises at the level of types, not tokens. Fred embodies a different type of point of view than typical humans. Jackson's point in the quoted selection might be restated as follows. One might say that although we cannot know what it is like to *be* Fred, we can, after surgery, know what it is *like* to be Fred. The possibility of this latter kind of knowledge is what Jackson sees as incompatible with physicalism; the impossibility of the former kind of knowledge stands regardless of the truth or falsity of physicalism.

It will be useful to consider a schematization of the knowledge argument. Since the version featuring the thought experiment about Mary is the one most frequently discussed in the literature, I will schematize the argument in terms of that one.

- First premise: prior to having a red experience, Mary can know all of the physical facts.
- *Second premise*: upon having a red experience for the first time, Mary comes to know for the first time what it is like to see red.
- *Third premise*: knowledge of what it is like to see red is knowledge of a fact.
- *Conclusion*: since Mary knew all of the physical facts prior to seeing red, the new fact she comes to know upon seeing red is a non-physical fact.

The conclusion of the knowledge argument constitutes a falsification of physicalism where physicalism is construed as the thesis that all facts are physical facts. Physicalists have been eager to find flaws in the knowledge argument. I briefly remark on some of the most popular kinds of physicalist responses to the knowledge argument.

Most physicalists responding to the knowledge argument have granted the argument's validity while questioning its soundness. The major kinds of physicalist responses may be grouped by which of the premises they question. One premise that has been subjected to quite a bit of physicalist pressure is the third premise stating that knowledge of what it is like is knowledge of a fact: propositional knowledge or knowing-that. Nemirov (1990) and Lewis (1990) have defended the "ability hypothesis": the hypothesis that knowledge of what it is like is constituted not by knowing-that but instead knowing-how, procedural as opposed to propositional knowledge. The abilities that Mary gains in learning what it is like to see red are the abilities to remember seeing red and to imagine seeing red. Intuitively, these abilities are the sort that could only be acquired by first having the experience of seeing red, thus explaining why prior to seeing red, physically omniscient Mary did not know what it was like. The ability hypothesis itself has come under attack. The strongest objections to the ability hypothesis hinge on pointing out that these abilities are not necessary for knowing what it is like (Alter, 1998; Tye, 2000). Tye makes the point especially forcefully by pointing out the degree to which the specificity of experience typically outstrips our abilities to remember or imaginatively relive those experiences. We may have an experience of a particular shade of red, call it "red 35," but be incapable of reidentifying precisely that shade or calling to mind precisely that shade and not some other. Nonetheless, at the moment in which we are experiencing that shade of red, we know, at that moment, what it is like. Thus knowing what it is like cannot consist in the abilities mentioned in the ability hypothesis.

Such criticisms of the ability hypothesis highlight the appeal, even to physicalists, of viewing knowledge of what it is like as a kind of propositional knowledge. In keeping with this line of thought is perhaps the most popular kind of physicalist response to the knowledge argument: the mode of presentation response. According to the mode of presentation response, all of the premises of the knowledge argument are true, but the argument itself is invalid. The key point is that even though there is a sense in which Mary gains new propositional knowledge, it is only in the sense in which she comes to represent an old fact in a new way—she comes to subsume an old fact under a new mode of presentation. The acquisition of Mary's new knowledge is akin to the transition from knowing that Mark Twain wrote *Huckleberry Finn* without knowing that Mark Twain is one and the same as Samuel Clemmons to knowing both that Mark Twain wrote *Huckleberry Finn* and that Mark Twain is Samuel Clemmons. In learning that Samuel Clemmons is the author of that book one comes to represent in a new way a fact that one knew already: the fact that Mark Twain wrote that book.

Questions arise, however, of what it is that Mary knew all along, what the new mode of presentation is, and what subjectivity amounts to on this view. For an especially clear answer to these questions, I turn to Lycan's version of the mode of presentation response. First off, note that Lycan endorses a representational theory

of conscious experience. 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. Mary is able to think to herself correctly the thought she would express by saying “*this* is what it is like to experience red.” Prior to having the experience, she was not in a position to employ correctly the indexical/demonstrative thought. But the state of affairs she picks out with the demonstrative is one she was able to pick out by other means prior to having the experience. This is analogous to how, prior to visiting my house, you are incapable of correctly thinking to yourself the thought you would express by saying: “Pete Mandik lives *here*.” Though you cannot refer to my home with “here” until you get here, you can nonetheless refer to it by other means prior to your arrival, e.g. under the description “Pete Mandik lives at *such-and-such* address.” Thus, according to Lycan, indexicality is the ultimate explication of subjectivity—as Lycan puts it, subjectivity is “pronominal.” Other accounts of subjectivity make similar appeals to the role of indexicals in thought—see, for instance, Tye (1995) and Rey (1997). But the insistence that subjectivity is indexical is developed at greatest length by Lycan.

One of the main virtues of the indexical response is the way that it connects with core concepts in standard articulations of the notion of subjectivity: concepts of perspective and point of view. Subjectivity is often identified with what is known from the first-person point of view where objectivity is associated with the third-person point of view. The relation of subjectivity to indexicality is evident when we turn to the role played by notions of point of view in the description of written language. Prose written from the first-person point of view is peppered with 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.” In contrast, typical instances of prose written from the third-person point of view are devoid of indexicals, as in “Herman Melville wrote *Moby Dick*.”

Despite such advantages, the indexical response is deeply flawed. One of the main problems is one that it faces simply in virtue of being a version of the mode of presentation response. Recall that prior to experiencing red, Mary is alleged to know *all* of the physical facts. This would include, among other things, facts about modes of presentation. Mary knows that “Mark Twain” and “Samuel Clemmons” are different modes of presentation of the same person, and “Hesperus” and “Phosphorous” are different modes of presentation of the same planet. Thus, in

learning something new in experiencing red, Mary cannot be learning a new mode of presentation of an old fact. (Similar objections are raised by Jackson, 1997, and Alter, 1998).

Besides this blanket objection to mode of presentation responses, there are objections that apply specifically to the indexical response. One such 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 Jackson puts a related point:

[T]he knowledge Mary lacked which is of particular point for the knowledge argument against physicalism is *knowledge about the experiences of others*, not about her own ... The trouble for physicalism is that, after Mary sees her first ripe tomato, she will realize how impoverished her conception of the mental life of others has been *all* along. She will realize that there was, all the time she was carrying out her laborious investigations into the neurophysiologies of others and into the functional roles of their internal states, something about these people she was quite unaware of. (1995, p. 181)

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. Perhaps an analogy can be helpful here. Questions of *what it is like to be X* are questions cast at the level of types, not tokens—likewise for questions of *what X looks like*. You probably know what Da Vinci's painting, *The Mona Lisa*, looks like, even though you probably have not seen the original. Suppose that I have hanging in my apartment a painting I did in high school called *A Study in Originality*, which is a color-by-numbers replica of *The Mona Lisa*. Given that you know what *The Mona Lisa* looks like, and that *A Study in Originality* looks like *The Mona Lisa*, you thereby know what *A Study in Originality* looks like. And all of this is consistent with the fact that you have never been to my apartment, and have not yet seen the particular painting *A Study in Originality*. You may not even be sure that I have any paintings in my apartment: you may not know *whether* I have any paintings. Nonetheless, *if* I have a replica of the *Mona Lisa*, and you know what the *Mona Lisa* looks like, then you know what the replica looks like: you know what it *would* look like were there

one, and you need not have contact with that particular instance of the type, *paintings-that-look-like-the-Mona-Lisa*.

These flaws with the indexical account of the subjectivity of consciousness motivate the pursuit of an alternative, which I provide in the following section. Before continuing, I preview the kind of response to the knowledge argument that will grow out of this account of subjectivity. First I list here what I think is right about the knowledge argument. I think that it is right to say that Mary would learn something new in having a red experience for the first time—prior to having a red experience, Mary could not know what it is like to see red. Further I think that knowledge of what it is like is a kind of propositional knowledge. But, in rejecting the mode of presentation response, I reject the claim that Mary is merely representing an old fact in a new way. What is wrong with the knowledge argument, I shall argue, is the first premise—the premise that states that prior to experiencing red Mary is able to know all of the physical facts. I think, instead, that Mary is able to know only all of the *objective* physical facts. The sorts of facts that she is cut off from are *subjective* yet nonetheless wholly physical facts. This sort of response has been suggested by others (Alter, 1998; Deutsch, undated; Flanagan, 1992), but the credibility of the suggestion suffers from the meager detail heretofore provided about what the physical basis of subjectivity amounts to. I turn now to provide such details, beginning with a reexamination of the notions of perspective and point of view.

2. Mental representation and pictorial perspective

The indexical account of subjectivity was on the right track insofar as it related the notions of perspective and point of view to representations. But it runs afoul of problems such as tying the subjectivity of experience to tokens instead of types. The indexical theorist analogizes mental points of view to language written and spoken in the first-person point of view. There is a less literary and more literal way of understanding perspective and point of view. This is the rather literal sense that 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 representational contents of the photographs produced include content about these points of view. This is why we can tell by looking at the photographs whether the camera occupied a point of view in front of or to the side of the subject.

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




FIG. 1: (a) and (b) are drawn in vanishing-point perspective, thus positioning the viewer in front of the first and second houses, respectively; (c) is drawn in oblique perspective and thus does not specify a point of view in front of any particular house.

rather than below it. 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 1).

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 1a and 1b are drawn in vanishing-point perspective. The vanishing-point in Figure 1a is located at the peak of the first house from your left, and in Figure 1b, 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 1c 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 1a specifies a point of view in front of the first house and Figure 1b specifies a point of view in front of the second house. Figure 1c, drawn in oblique perspective, abstracts away from which of the three houses the viewer is in front of, since, unlike in Figures 1a and 1b, the sides of the three houses occupy equal portions of the picture plane. Nonetheless, Figure 1c 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 by considering the notion of 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 1a 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 (see e.g. 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., Bühlhoff & 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. O'Keefe & 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 the hippocampus that they investigated were allocentric representations of spatial location. Where the 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. (See Stein, 1992, and Milner & Goodale, 1995, for general discussions of egocentric representations in parietal cortex.) However, the question of *where* egocentric representations are is not as interesting as the question of *what* egocentric representations are. I propose to explore such questions against the backdrop of naturalistic theories of representational content.

Here I will assume an informational psychosemantics, the crux of which is the view that a brain state represents X by virtue of having the function of carrying information about (being caused by) X . While informational psychosemantics offers these conditions as both necessary and sufficient for representations, here I need only to regard them as sufficient. Elsewhere I argue against their necessity (Mandik, 1999) and offer as an additional set of sufficient conditions that X may represent Y if X has the function of causing Y . As discussed there, this (procedural instead of informational) account is most appropriate to accommodating the view, widespread in neuroscience, that activations 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 will 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 (affecter) 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 paradigms for causal covariational semantics are feature detectors. A feature detector consists in 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. Neuroscientists have hypothesized the existence of detectors for stimulus types such as high-contrast edges, motion direction, and colors. Some have hypothesized that the nervous systems of frogs contain fly detectors. 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).

Other examples of feature detectors include 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 causal 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 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 through Z_n that 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 neurons, whether in cortex or in the retina itself, demonstrate patterns 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: $_ \wedge \wedge \wedge _$. 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 offline utilization of perceptual (and perhaps motor) processes that are typically used online (see, for example, Grush, 2001). 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 representation 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 covariational 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 carries 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 offline what is run online during perception: states that are supposed to carry information in the perceptual case may also be pressed into service for offline imaginings. Thus the offline states employed in imagining the Eiffel Tower owe their representational content to the information they are supposed to carry in the online 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 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 the 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 a 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) piggybacks 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. Another way to put the point is that when pictures are used indexically to represent a particular object, as in the case of a photograph of a particular car, they can implicate a *particular* point of view. But pictures can also be put to a use whereby no particulars are represented and no particular points of view are implicated. In the case of a drawing of a general *type* of car, the point of view implicated is likewise a *type* of 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.

Thus completes my sketch of pictorial perspective and its neural underpinnings. All the examples I employed involved spatial representations. One obstacle to claiming that pictorial perspective is an appropriate way of thinking of the subjectivity of conscious experience will depend on extending the account to non-spatial cases. While we experience spatial locations and relations, we also experience non-spatial features of the world, including colors, odors, and temperatures [1]. In order to show that the account of perspective I favor can be extended to non-spatial cases, I next discuss how experiences of temperature may involve perspectival representations of temperature.

In a detailed examination of thermoperception (drawing on the work of Hensel, 1982), Akins (2001) 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, 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 a well-known 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?” She takes the narcissism of thermoperception to count against the claim that thermoreceptive 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. 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, for example, 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. Thermoreceptors 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: representations of chemical-involving 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*" (Akins, 2001, p. 383, emphasis in original). 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 thermoreception 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 [2].

I have proposed that some mental representations exhibit pictorial perspective. We need to tie this into consciousness. Do states of consciousness possess this kind of perspective? Regarding whether conscious states exhibit this kind of perspective, the answer seems a resounding "yes." The thermoreception 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.

This completes my sketch of the pictorial account of how mental representations may be perspectival. I turn now to address the problems about perspective and subjectivity raised by the knowledge argument.

3. Pictorial perspective and the knowledge argument

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. 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 an image of *X* represents is the way *X* would look like from some particular location. What the image represents depends for its existence on the process of its representation in the sense that 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. Consider a pictorial representation, such as a photograph of a complex object like the Statue of Liberty. Given the point of view from which the photograph was taken, only a fraction of the surface of the statue is explicitly represented in the photograph. Certain regions seen from one angle would be obscured from another angle. Consider the set of regions that are captured by the photograph: the set of all and only the regions explicitly represented in the photograph. 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 they would exist even if no one were to represent them. But in specifying the set comprised of all and only the spatial regions captured in the image, one does not carve nature at the 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 a particular type of 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 of 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. There are properties represented in sensory experience that may only be represented in sensory experience. *Part* of what is represented in olfaction, say, may be conveyed in some other mode of representation like spoken language—“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.

On my view, restricted epistemic access to some fact is restricted representability. Something that can be known in only a restricted number of ways is something that can be represented in only a restricted number of ways. Conversely, something that is knowable in a multiple ways can be represented in multiple ways. The causal covariational psychosemantics sketched above helps to spell out this notion of restricted epistemic access. 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 property P is knowable from multiple kinds of point of view, then there are multiple kinds of representations— R_1 through R_n that may represent P and thus multiple kinds of causal interactions that P may enter into. Suppose instead that property P enters into only a restricted number of causal interactions. This is indeed a physical possibility. Nothing in physicalism rules it out. Such a property, in virtue of entering into only a few kinds of causal interactions is thus representable, if at all, in only a few number of ways. What might examples of such properties and interactions be? Consider the old torn Jell-O box trick. 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-forgable identifier. The matching halves are so informationally complex that the only way of identifying one is by identifying the other. As Dennett puts it,

[t]he only *readily available* way of saying what property M is [the property that makes one piece of the Jell-O box a match for the other] is just to point to our M detector and say that M is the shape property detected by this thing here. (p. 637)

As I would state the case, the only *detection supporting* causal interactions that one piece of the Jell-O box enters into are causal interactions involving the other piece of the box. In the psychosemantics glossed above, the paradigm examples of representation are sensory detection, and the example of the Jell-O box helps us appreciate the possibility that many of the properties represented in sensory mechanisms are properties that enter into only a few *detection supporting* causal interactions,

namely, interactions with the senses. And someone without those sensory mechanisms could not represent those physical properties.

What, then, did the color-blind Mary not know? There was a physical property that she was incapable of representing until she first had her red experience, because this physical property enters into detection supporting causal interactions only with the mechanisms of human color vision. Note, however, that it is knowledge about a *property*, not knowledge restricted to the particularity of her own case. Thus, unlike the indexical account criticized above, the present account does not mistakenly tie subjectivity to tokens but to types.

What more can be said about the property that Mary comes to have knowledge of? By definition, not much. Such a property, if representable only by certain sensory experiences, by definition could not be represented by linguistic representations. Besides telling you “it’s *that* property”—which is obviously uninformative—nothing more can be said since this subjective property, this subjective fact, does not enter into the rich number of kinds of detection supporting causal interactions that objective properties and facts enter into.

The possibility sketched above is one no physicalist should have qualms about admitting. Physicalism provides no reasons to expect that all physical properties be objective in the above senses. Physicalism does not rule out the possibility of a property that is instantiated in causal remoteness from most other properties and is thus represented, if at all, only in a restricted number of ways. Nonetheless, familiarity with the sorts of considerations invoked in Nagel and Jackson’s thought experiments may incline us to contrary intuitions. Students of these thought experiments are accustomed to imagining (or at least pretending to imagine) that a suitably informed human mind could be wrapped around the entirety of the objective physical facts about some domain. For instance, it is oft supposed that *all the physical facts* about bats could be known, or at least represented, from the objective point of view without having batty experiences. But it is precisely this kind of supposition that I call into question.

Much of the problem, I want to suggest, has to do with ambiguities in our word “represent.” There is one sense, I will grant, that we can represent all of the physical properties from the objective point of view, and another sense in which we cannot. And it is this latter sense that is crucial, I suggest, in the representations supporting knowledge of *what it is like*. Consider the linguistic representation “All of the people in France are less than 3 meters tall” and the fact that my friend, Jacques Blanc, is in France and is less than 3 meters tall. Does the linguistic representation in question represent the fact in question, the specific fact about Jacques Blanc? I take it as obvious that there is a sense in which it does and a sense in which it does not. The sense in which it does is the sense in which it says something true of Jacques Blanc. The sense in which it does not is the sense in which it is not specifically about Jacques—it does not single out Jacques from the other inhabitants of France. Sensory experience is analogously specific and so is the accompanying knowledge of what it is like to have specific kinds of experiences. One who merely knows what it is like to see colors may not know what is like to see some specific color. One who has seen only shades of blue and yellow does not thereby know what it is like to see

red, just as the normal humans did not know what it was like to see Fred's red₁ and red₂ in Jackson's thought experiment. The properties delivered in sensory experience are quite specific. But the ways that we represent physical properties in thought experiments asserting the possibility of physical omniscience from the objective point of view are devoid of the requisite specificity. We imagine a bunch of electrons here and a bunch of protons there, a bunch of neurons here and a bunch of ganglia there. Physical theory provides us with the means of the relatively unspecific linguistic representation "everything is made of electrons, protons, etc." But physical theory does not provide us the means of identifying one half of the Jell-O box without recourse to the other half. Nor should we expect it to. Further, and more importantly, this gives us no ground for concluding that Jell-O box halves are non-physical. Likewise physical theory does not provide us with the means of representing all and only the properties represented in bat experience. And we should neither expect it to nor conclude from this failure that there is something non-physical that we are thereby cut off from.

4. Conclusion

The pictorial account of perspectival mental representation allows us to see how subjectivity may be physical in a way that squelches the threat of the knowledge argument. I have noted several differences between the present account of subjectivity and previous accounts. For instance, I have emphasized the way in which subjectivity comes in at the level of types, not tokens. I close by noting a further difference not yet made explicit. Accounts of subjectivity like those provided by Lycan (1996) and Tye (1995, 2000) portray subjectivity as entering the picture with the introspective application of concepts. On their account, what makes knowledge of what it is like perspectival is the *way* the thing is known, not the *thing* that is known. On such views then, if there were no creatures capable of applying second-order representations to their first-order experiences, there would be no subjectivity. On the view I advocate, subjectivity enters the picture with the first-order states. This is neither the time nor the place to attack higher-order theories of consciousness like Rosenthal's (1986), but *if* there is something it is like to be a creature that has first-order experiential representations even if that creature has no second-order representations, *then* there is something it is like to be the creature whether the creature knows it or not. There may be something it is like to be a bat whether bats themselves know it or not. Regardless of whether bats know it or not, *we* do not know what it is like to be a bat. Why not? Because facts of what it is like are subjective physical facts, facts that depend on experiences, even though they may not depend on higher-order representations of those experiences.

Acknowledgements

For especially detailed and useful comments on oral presentations of this material I thank Robert van Gulick and Mark McCulloch. I am also thankful for comments on previous versions of this material from Bill Bechtel, Peter Machammer, Rick Grush,

Tad Zawidzki, Jonathan Waskan, Eric Steinhart, Chase Wrenn, Andy Clark, William Lycan, Jesse Prinz, Robert Gordon, Murat Aydede, Wade Savage, David Rosenthal, Brian Keeley, Chris Eliasmith, John Bickle and Christopher Gauker. For useful feedback I thank also audiences of presented versions of this material at meetings of the Society for Philosophy and Psychology, the Eastern Division of the American Philosophical Association, the New Jersey Regional Philosophical Association, the McDonnell Project for Philosophy and the Neurosciences, the Philosophy Department of Washington University in St. Louis, the Philosophy Department of William Paterson University of New Jersey, and the Theoretical Cognition Group at the University of Pittsburgh.

Notes

- [1] For further discussion of spatial experience and its connection to the topics of objectivity and subjectivity, see Mandik (1998, 1999).
- [2] For additional discussion of Akins' account of thermoperception, see Bickle and Mandik (1999).

References

- AKINS, K. (2001). Of sensory systems and the "aboutness" of mental states. In W. BECHTEL, P. MANDIK, J. MUNDALE & R. STUFFLEBEAM (Eds) *Philosophy and the neurosciences: a reader*. Oxford: Blackwell.
- ALTER, T. (1998). A limited defense of the knowledge argument. *Philosophical Studies*, 90, 35–56.
- BICKLE, J. & MANDIK, P. (1999). The philosophy of neuroscience. In *The Stanford encyclopedia of philosophy*. Available online at: <<http://plato.stanford.edu/entries/neuroscience>> .
- BÜLTHOFF, H.H. & 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.
- DENNETT, D. (1997). Quining qualia. In N. BLOCK, O. FLANAGAN & G. GÜZELDERE (Eds) *The nature of consciousness*. Cambridge, MA: MIT Press.
- DEUTSCH, M. (undated). Subjective physical facts. Unpublished.
- EICHENBAUM, H., STEWART, C. & MORRIS, R.G.M. (1990). Hippocampal representation in spatial learning. *Journal of Neuroscience*, 10, 331–339.
- FEIGENBAUM, J. & ROLLS, E. (1991). Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology*, 19, 21–40.
- FLANAGAN, O. (1992). *Consciousness reconsidered*. Cambridge, MA: MIT Press.
- GRUSH, R. (2001). The architecture of representation. In W. BECHTEL, P. MANDIK, J. MUNDALE & R. STUFFLEBEAM (Eds) *Philosophy and the neurosciences: a reader*. Oxford: Blackwell.
- HENSEL, H. (1982). *Thermal sensations and thermoreceptors in man*. Springfield, IL: Charles Thomas.
- JACKSON, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–136.
- JACKSON, F. (1995). What Mary didn't know. In J.D. TROUT & P.K. MOSER (Eds) *Contemporary materialism: a reader*. New York: Routledge.
- JACKSON, F. (1997). Three questions about consciousness: review of *Consciousness and experience* by William G. Lycan. *Psyche*, 3(5). Available online at: <<http://psyche.cs.monash.edu.au/v3/psyche-3-05-jackson.html>> .
- KOSSLYN, S.M. (1994). *Image and brain*. Cambridge, MA: MIT Press.
- LETTVIN, J.Y., MATURANA, H.R., McCULLOCH, W.S. & PITTS, W.H. (1959). What the frog's eye tells the frog's brain. *Proceedings of the IRF*, 47(11), 1940–51.
- LEWIS, D. (1990). What experience teaches us. In W. LYCAN (Ed.) *Mind and cognition: a reader*. Oxford: Blackwell.
- LYCAN, W. (1996). *Consciousness and experience*. Cambridge, MA: MIT Press.

- MANDIK, P. (1998). Objectivity without space. In R. GRUSH (Ed.) *The Philosophy of Gareth Evans: Special Issue of the Electronic Journal of Analytic Philosophy*. Available online at: <<http://www.phil.indiana.edu/ejap/ejap.html>>.
- MANDIK, P. (1999). Qualia, space, and control. *Philosophical Psychology*, 12, 47–60.
- MILNER, A.D. & GOODALE, M.A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- NAGEL, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.
- NEMIROW, L. (1990). Physicalism and the cognitive role of acquaintance. In W. LYCAN (Ed.) *Mind and cognition: a reader*. Oxford: Blackwell.
- O'KEEFE, J. & NADEL, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- REY, G. (1997). *Contemporary philosophy of mind*. Oxford: Blackwell.
- ROSENTHAL, D. (1986). Two concepts of consciousness. *Philosophical Studies*, 49, 329–359.
- 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*, 15, 691–700.
- TYE, M. (1995). *Ten problems of consciousness: a representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.
- TYE, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
- ULLMAN, S. (1989). Aligning pictorial descriptions: an approach to object recognition. *Cognition*, 32, 193–254.