In *Perception, Vancouver Studies in Cogntive Science, Vol. 5*, ed. Kathleen Akins, Oxford: Oxford University Press, 1996.

2

Explaining Why Things Look the Way They Do

Kirk A. Ludwig

1. Introduction

How are we able to perceive the world veridically? If we ask this question as a part of the scientific investigation of perception, then we are not asking for a transcendental guarantee that our perceptions are by and large veridical: we presuppose that they are. Unless we assumed that we perceived the world for the most part veridically, we would not be in a position to investigate our perceptual abilities empirically. We are interested, then, not in how it is possible in general for us to perceive the world veridically, but instead in what the relation is between our environment and its properties, of which we have knowledge, on the one hand, and our perceptual mechanisms, on the other, that results in very many, even most of our perceptions being veridical in everyday life.

In this paper, I am concerned with a certain kind of answer to our question which has been popular in psychological studies of our perceptual abilities at least since Helmholtz (1867).[1] The answer is that we do it by taking account unconsciously of various perceptual cues about objects and events in our environment and then reasoning to what the environment must be like on the basis of these cues, our general knowledge of the environment, and how it impinges on our perceptual organs. It is doubtful that anyone has ever held a pure inference theory. For present purposes, I will call any theory an inference theory that appeals at least in part to unconscious inferences from cues provided by stimulus to the nature of the perceiver's environment in explaining how things look. While the power of these accounts is undeniable, they are, I think, deeply mistaken. When I say this, I do not mean merely that they are as a matter of fact false, or that the evidence in fact is overwhelmingly against them. I mean that the explanans employed *could not* explain the explanandum; the appearance of explanatory force is an illusion. It is not that all the particular

explanations of this type have failed, but that no explanation of this type could be correct.

In what follows, I will first give some examples of the sort of explanation that is the target of this investigation. I will concentrate on two examples drawn from explanations of visual perception, the visual perception of size and motion,[2] but many of the points I will make will generalize to other sensory modalities, and to explanations of non-perceptual cognitive capacities as well.[3] My initial aim will be to present these explanations in as strong a light as possible. I will then develop an *a priori* argument to show that they cannot be correct, develop some detailed criticisms of them, and provide a diagnosis of their appeal. I will, however, also argue that despite these explanations being necessarily false, they can be reinterpreted so that they have a legitimate use in psychological investigations of perception. The empirical evidence that psychologists have accumulated is not unimportant, but it is not evidence for the existence of unconscious thoughts and inferences. I will conclude by replying to some possible objections to my argument.

2. Perceptual Achievement as the Result of Unconscious Inference: The Visual Perception of Size and Motion

We take our environment to be a three-dimensional space filled with objects, sound and light which evolves continuously through time. We conceive of ourselves as located within that space, as capable of moving around in it, and as subject to stimulus from surrounding objects. Our problem is to explain how we recover from the stimulus an accurate representation of our environment at a time and through time.

The objects around us are located at different distances and directions from us, are in motion or at rest, and their state of motion and distance from us affect the kind of stimuli we receive from them, if any. To achieve an accurate perceptual representation of their properties, we need to take account both of the similarities among them and their differences. Since their locations make a difference to the stimulus which we receive from them, one of the perceptual tasks we must solve is how to represent as the same those properties of objects (or parts of objects) from which we receive differing stimuli solely because of their different locations. This is the problem of perceptually representing constancies across differences in location. The general solution to this problem for visual perception is to provide a function that maps the physical stimulus on the retina to a representation of the environment which meets the constraint that representations so generated are by

and large veridical. This function may take into account more than just the local stimulus on the retina.

Perceptual experience in general is autonomous from cognitive states such as beliefs, judgements, assumptions, suppositions, hypotheses, etc. This is what allows our perceptual experiences to serve as evidence for our judgements about the world around us. The autonomy of experiences from cognitive states means that how a scene perceptually appears to us, how a scene looks, is, by and large, independent of what we believe about it, suppose about it, judge about it, etc. This is most clear in the case of perceptual illusions or hallucinations which we know or believe to be such. It may visually appear to me that there is a pile of leaves on my bed while at the same time I do not believe this because I believe that I am undergoing a drug-induced hallucination. Similarly, as in the Müller-Lyer illusion (Figure 2.1a), one line on a page may appear to me to be longer than another, although I know after measuring them that they are of equal length. The autonomy of perception is even more dramatically illustrated in cases of impossible figures, such as the Penrose triangle in Figure 2.1b. Thus, it is clear that my belief that my experience is non-veridical need not change the character of my visual experience, for if it did, there could not be such a thing as undergoing a hallucination, or experiencing an illusion, while simultaneously believing it to be one. The representational constancies we are interested in, then, are not to be found in what we believe about our environment, but in our perceptual representations of them.

As a general term to cover representational features of our perceptual experiences, as opposed to our beliefs about and other attitudes toward those experiences, or what we believe about the world on their basis, we can use the expressions 'how things look' and 'how things

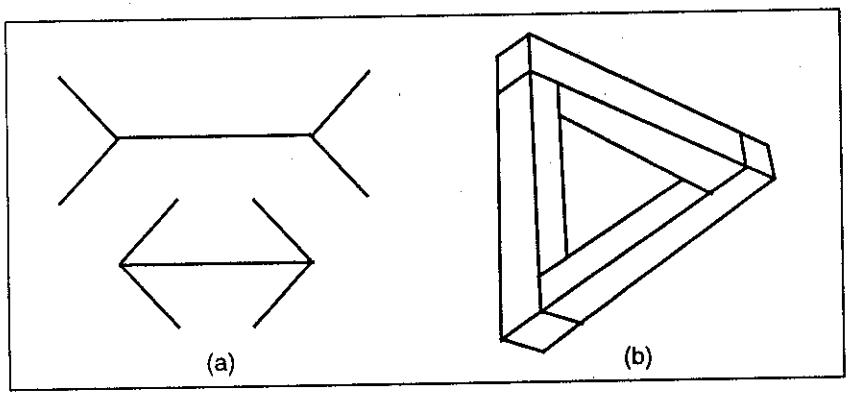

Figure 2.1: (a) The Müller-Lyer Illusion. (b) The Penrose Triangle.

appear', and their variants. Thus, we might speak, for example, of how long a line looks, or how light a thing appears to be, to designate a representational feature of a perceptual experience. Ultimately, we want to talk about the representational character of a visual experience independently of any claim about the objects, if any, of the experience, so that we can talk of the representational character even of total hallucinations; we can use the expressions we introduced above in this sense, that is, to talk about the representational character of an experience from the point of view of the person whose experience it is, without committing ourselves to there being any thing or things the person's experience is of.

Veridical Perception of Size

The correct perceptual representation of size constancy requires that an object's size appear the same when it changes its location with respect to us, and that objects at different locations appear to be the same size if they are. The inference theory offers a compelling account of how the perceptual system accomplishes this task. The perceptual system must recover from the stimuli impinging on the individual, in this case, from the light falling on the retina, a representation of the (relative) sizes of objects in the environment. The actual stimuli from an object relevant to its size is the area occluded by its image on the retina, or, as we will say, its visual angle. (To keep the issues in clear focus, we will set aside changes of visual angle due to changes in an object's orientation, which would have to be treated in a fully general account.) An object's size is proportional to the distance it is from the perceiver and the visual angle of its image on the retina. Thus, the perceptual system can generate an accurate representation of the size of an object provided that it first determines and keeps track of the distance of the object from the observer. Therefore, if we assume that the perceptual system has determined from other cues the distance of an object from the observer, we can explain how the perceptual system keeps track of the size of an object in a visual representation of it by postulating that it infers from the distance of the object and its visual angle how large it is. For example, if the visual angle increases while the distance decreases proportionally, the visual system infers that the object's size has not changed. This explains why an object that approaches us does not look as if it is getting larger. This process is represented in Emmert's law: perceived distance × visual angle = perceived size.[4]

We can contrast the inference theory's explanation of size constancy with that of the stimulus theory (see, e.g., Gibson 1950). The stimulus

theory aims to explain size constancy in experience without appeal to unconscious inference. Thus, the dispute between the inference theory and the stimulus theory has the character of an empirical dispute between two contingent theories of the same phenomenon. According to the stimulus theory, size constancy is explained not by unconscious inferences but instead in terms of a constant property of the sensory input which is directly correlated with the representation of an object's size, e.g., constant ratios of the visual angle of an object to the visual angle of some appropriate frame. Thus, in Figure 2.2, *A* is perceived to be the same size as *B* because the number of units of the grid occluded by *A* is the same as the number occluded by *B*. Experiments have demonstrated that ratios of this sort have an important role to play in how things look to us (Rock and Ebenholtz 1959).

However, although constancies in the ratios of the visual angles of objects to a frame seem to be a factor in the perception of size constancy, the stimulus theory does not accurately predict the extent of the appearance of size constancy, and is not applicable to all situations in which size constancy is achieved. For example, the *apparent* size of a line in a rectangular frame, which has the same ratio to the frame as a nearer line and frame, is smaller than that predicted by the stimulus theory. In addition, size constancy can be attained by subjects in a dark room when viewing a single luminous object as long as distance information is available. This appears to be a strong argument, then, in favour of the inference theory providing at least an important part of the account of how size constancy is achieved in visual perception.

Further support for the inference theory is provided by its ability to explain systematic failures to achieve veridical perception. Consider the familiar Ponzo illusion, e.g., illustrated in Figure 2.3. In this illu-

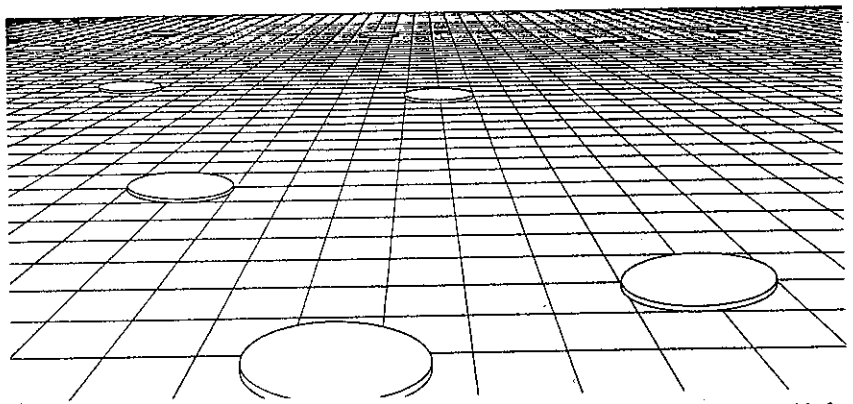

Figure 2.2: The Ratio Theory of Size Constancy. From Irvin Rock, *Perception* (New York: W. H. Freeman, 1984), p. 338. Used with permission.

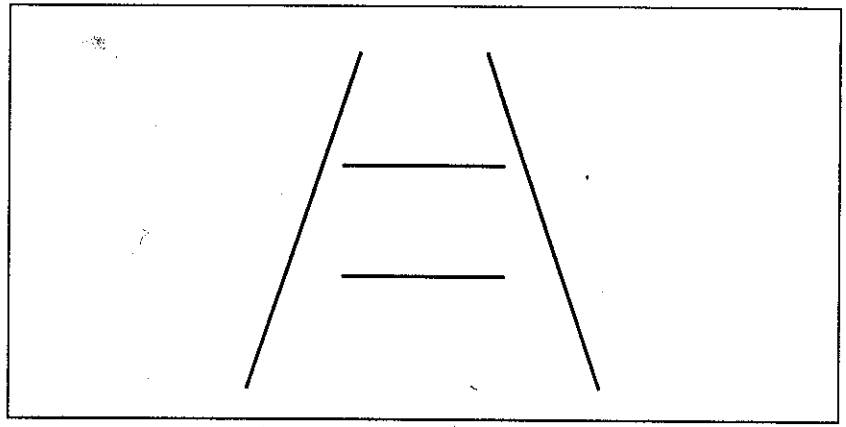

Figure 2.3: The Ponzo Illusion

sion, the upper line appears to be larger than the lower line, although they have the same length. The inference theory explains this as the result of a mistaken inference. The two lines which converge toward the top of the figure are interpreted by the visual system as indicating increasing depth, because the image projected on the retina is similar to the image which would be projected on the retina by two parallel lines receding in a horizontal plane under the observer and orthogonal to him. Thus, the upper line is inferred to be at a greater distance than the lower line, but it produces a visual angle that is equal to that produced by the lower line. In accordance with Emmert's law, the visual system infers that the upper line is larger.

This gives us three sources of support from the inference theory of size constancy. The theory accounts for the veridical perception of relative sizes. It provides an account which at least in some cases is better than its main rival in this area, the stimulus theory. And it explains not just veridical perception of size constancy, but also systematic breakdowns in veridical perception of relative size.

Veridical Perception of Motion

A natural first suggestion for how the perceptual system keeps track of motion is that an object is perceived to move provided that the image it projects onto the retina changes its position on the retina, and that its rate of motion is proportional to the rate of motion of its image on the retina. This would be a pure stimulus theory of the visual perception of motion. A moment's reflection shows that this theory cannot be the correct account of how we achieve veridical perception of motion, because the movement of an image on the retina is a function

not just of the movement of the object but also of the movement of the head, eyes and the body of the observer. Thus not every movement of the image on the retina is an indication that the object that projects it is moving. To perceive motion veridically the perceptual system must take into account the reason for the motion of the image on the retina in order to distinguish the motion which is due to the movement of the object from that due to the movement of the retina itself. Since we must appeal to a process that takes into account a variety of different sorts of information, we must appeal to an inference explanation of the veridical perception of motion, and not a simple stimulus theory. Thus, the perceptual system takes into account unconsciously the movement of the image on the retina, and then subtracts from it the movement due to the movement of the eyes and the head, and of the body through the environment, to arrive at a representation of the movement of the object projecting the image. If the eyes are stationary relative to the head, which is rotating to the left, the perceptual system will infer from an image moving at an equal rate in the same direction relative to the retina that the object projecting the image is motionless relative to the observer.[5]

Not every perception of motion depends just on information about the motion of the image relative to the motion of the retina. This is illustrated in the illusion that the moon is moving when it is seen through slowly moving clouds. The inference theory can explain this phenomenon in terms of the system's assumption that the clouds represent the background, and that the background in a visual scene is stationary relative to the observer. The only possible interpretation of the relative motion of the image, then, is that the moon is itself moving. The value of using relative motion of images on the retina as a source of information is that it is sensitive to small changes of position which are otherwise difficult to detect.

As in the case of the visual perception of size constancy, some striking evidence for the inference theory is provided by its ability to explain why we perceive what we do in cases of non-veridical perception of motion, as in cases of stroboscopic or phi phenomena. This is also a case in which the inference theory is apparently able to provide a better explanation than the stimulus theory. Consider two shapes flashed alternatively on a screen, as in Figure 2.4a. It has long been known that at the right speed of alternation the dot will appear to move from position A to position B. How is this to be explained? This presents a difficulty for a simple stimulus theory according to which movement is perceived provided that an image moves across the retina, for in this case there is no movement of an image across the retina at all. In contrast, the inference theory can explain this phenomenon

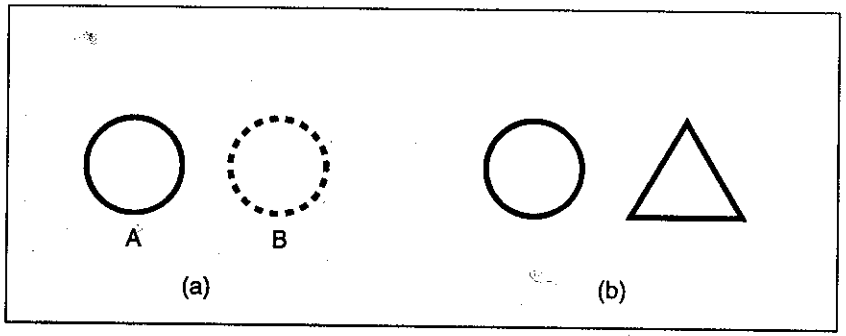

Figure 2.4: The Stroboscopic Effect

in the following way. In ordinary perception of objects which are moving rapidly back and forth, the movement of the object between the points at which it comes to rest can be rapid enough not to register on the visual system. In this instance, there is no image registered as moving across the retina. Instead, what is registered is the image at the end points of its movement. To perceive this as motion, the visual system has to take account of more than simply the information that is provided by the movement of an image across the retina. It must take account of the rapidity and sequence of the appearance and disappearance of the images. Thus, the explanation of the illusion of movement in the case of images flashed alternately on a screen at different positions in rapid succession is that the visual system infers from the alternation and speed of succession that an object is being moved back and forth rapidly from one location to another.[6]

The inference theory also explains the conditions under which the illusion vanishes: (i) when the alternation is slow enough that if an object were moving back and forth at a speed compatible with the alternating images, its image should be visible in between, and (ii) when the rate of alternation combined with the duration of the images is incompatible with an object accelerating and decelerating back and forth between two points.

An interesting example of this process is provided by alternating images of objects of different shapes, as shown, e.g., in Figure 2.4b. At the right speed of alternation, the circle is perceived both to move and to change its shape into a triangle, and then vice versa. Thus, the perceptual system appears to be inferring that the circle has moved because the rate of alternation is appropriate for rapid motion back and forth. But this presents a problem, since it implies that one object is moving, and the image at one end of the trajectory is of a different shape than at the other. The perceptual system solves the problem by inferring that the shape of the object is changing as well, and so representing it.

Further evidence is provided by an experiment in which rectangles are moved over stationary spots on a screen, as illustrated in Figure 2.5 (Sigman and Rock 1974). The experiment is performed under two conditions. In the first, the rectangles are not visible to the subject. In this case, the spot appears to be moving back and forth. If the inference theory is correct, this is because the perceptual system has solved the problem of what the stimulus represents by hypothesizing that an object is moving back and forth rapidly. If so, then with additional information which provides an alternative explanation, the illusion of movement should disappear. This is precisely what happens when the two rectangles which move back and forth over the dots are visible to the subject. The spots are seen as stationary objects, each of which is alternately revealed and occluded. In this case, the perceptual system takes account of the additional information and rejects the hypothesis that the dots are moving back and forth rapidly. That the visible rectangles are occluding and revealing each dot successively provides a better explanation for the alternating images.

To show that the inference theory is genuinely empirical, that is, that it is falsifiable, we can consider one perceptual phenomenon involving apparent motion for which (it seems) it is clearly not the best explanation, the waterfall illusion. The waterfall illusion is generated by having a subject look steadily at moving contours over a uniform background, and then look at a set of stationary contours against a uniform background. The stationary contours appear to the subject

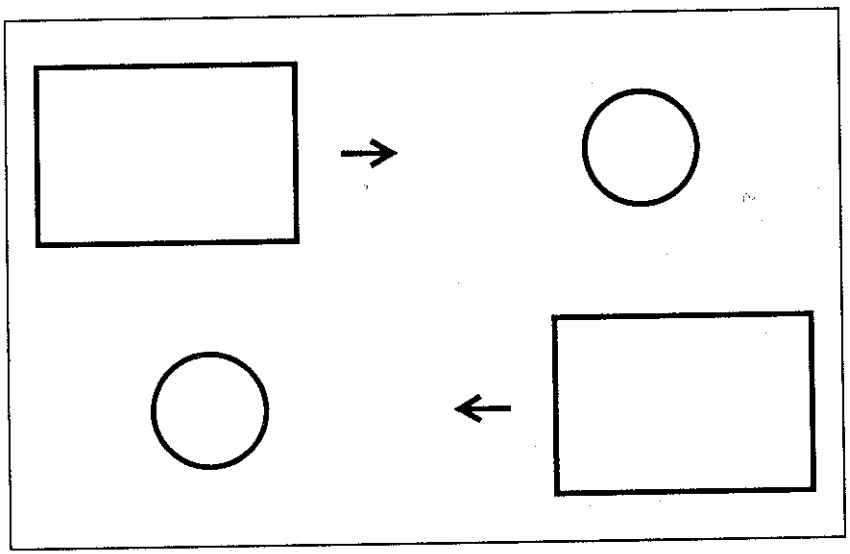

Figure 2.5: Destroying the Stroboscopic Effect

to be moving in the direction opposite to that of the moving contours she had previously been looking at. It does not seem that the inference theory can explain why this effect occurs, for there is nothing about either the image of the stationary contours or the succession of the image of the moving contours and then the stationary contours which would provide any reason to infer that the contours were moving in a direction opposite to that of the moving contours. In addition, there is a straightforward mechanistic explanation for this effect in terms of sensory adaptation to contours moving over the retina (Anstis and Gregory 1964). This occurs because some cells in the visual system are more sensitive to motion in one direct than another. When these become fatigued, cells sensitive to motion in the other direction are more active. In the absence of continued stimulus provided by image motion in one direction, the cells for detecting motion in the opposite direction now produce a relatively stronger signal, which produces for a short time an illusion of motion in the direction opposite to that of the motion most recently detected.[7]

There are two differences between this case and the cases in which the inference theory seems to provide a better explanation. The first is that in this case there is a simple mechanical explanation of the phenomenon which does not require taking into account various different sources of information to arrive at a best guess about the properties of the perceptual scene. The second is that in this case the perceptual effect is not in any plausible way thought of as the result of a misapplication of a process which in conjunction with the usual nature of the environment is likely to lead to veridical perception. In the Ponzo illusion, by contrast, the appearance is explained as a result of an inference which in the usual case produces a veridical perception. In the waterfall illusion, representing motion as the direction opposite to that most recently detected is not likely to lead to veridical perception in our normal environment. It seems to have no justification in terms of the goal of achieving veridical perception of one's environment.

Summary

In both of the applications of the inference theory we have considered so far the perceptual problem is at one level of description the same: it is to keep track of properties of objects located at different positions from us through changes in their positions and ours. In the first case, the property we want to keep track of is an object's relative size. In the second, it is an object's relative motion. In both cases, the inference theory seems to provide the best explanation because successful representation requires us or our perceptual system to take account of a

number of different sources of information simultaneously or sequentially in achieving veridical perception.

These are representative examples of how the inference theory explains veridical perception. The theory gains force as it is seen to be applicable to a wider range of perceptual phenomena. But that it can be applied naturally to a wide range of phenomena should not be surprising, for we can already see that the inference theory is most compelling when generating a veridical perception requires taking into account a number of different interrelated factors, and most of the properties in our environment which we want to keep track of generate peripheral stimuli as a function of their distance from us, ambient light, their orientation, intervening objects, their movement, our orientation, position, movement and so on. Thus, virtually any property we could wish to represent can be expected to admit of an explanation of this form.

In concluding this section, we can note a number of features of the inferences postulated by the inference theory which will play a crucial role in the criticism to come. First, the inferences postulated by the inference theory are clearly not conscious inferences or even readily accessible to consciousness. For example, in the Ponzo illusion, one is not aware of seeing the two lines as having the same visual angle, noting that the two converging lines can be interpreted as representing increased depth, and then consciously thinking that that means the upper line must be further away, and, hence, since it has the same visual angle as the lower line, larger. Instead, we simply see the upper line as larger. Nor does learning the theory help us to identify in ourselves these inferential processes. This is typical of the phenomena explained by the inference theory. We are neither aware unreflectively that any inference is taking place, nor are we able to bring it to our awareness by any act of attention or concentration.

Second, given the way evidence is marshalled for the inference theory, it is clear that the inferences postulated are not thought of as requiring accessibility to consciousness. The evidence for the existence of these inferences is exclusively third-person evidence. The warrant for postulating them derives from their providing the "best explanation" of our achieving veridical perception of the world around us. The absence of any first-person awareness of these inferences, that is, awareness by the subject of the inferences, or the perceiver, is not thought to be relevant to the confirmation of the theory. The absence of such accessibility, even in principle, is not thought to be a refutation of the theory. There is thus no logical requirement that these inferences be connected in any way with our conscious mental lives other than by their output being our visual experiences.

Third, there are positive reasons to require that these inferences be in principle inaccessible to our conscious mental lives. For what we want to explain are perceptual experiences whose contents are by and large autonomous with respect to our conscious mental lives or any of our mental states which dispositionally manifest themselves as conscious mental states. Appeal to such states could not help explain a perceptual phenomenon which we know is insensitive to what we can know, or believe, or want consciously. It is precisely because these states are both unconscious and inaccessible to consciousness that they can be used to explain perceptual experience. They must be as autonomous from our conscious mental lives as the experiences they are supposed to explain; otherwise they could not play the role they are designed to play in the explanation of perceptual experiences. To borrow Helmholtz's expression, unconscious inferences not only are but must be "irresistible."

3. The Connection Principle and the First-Person Perspective

Despite the powerful appeal of the inference theory, I will argue that it cannot be correct. It cannot be correct because it violates a central conceptual requirement on attributing to a person a mental state or process, namely, that that state or process be specially connected to that person's perspective on his own mental states, which is essentially connected to their manifestability in his conscious mental life. It is precisely the autonomy from our conscious mental lives of the inferences which the inference theory postulates which undermines the possibility of regarding them as genuinely mental states of the perceiver.

Searle's Argument

Following John Searle's recent discussion (1990a),[8] I will call this claim the *connection principle*, though my formulation of it and my argument for it are different from Searle's. Searle's version of the connection principle does not deny that there are unconscious mental states, but it requires that every unconscious mental state be potentially a conscious mental state. Searle gives several different formulations of the connection principle. At one place he puts it by saying that every mental state is potentially a conscious mental state, at another by saying that "the notion of an unconscious intentional state is the notion of a state that is a possible conscious thought or experience" (Searle 1990a, 588), and at another by saying that "that ontology of the unconscious consists in objective features of the brain capable of causing subjective

conscious thought" (Searle 1990a, 588). I don't think that any of these formulations are equivalent.[9]

I have argued elsewhere (1993) that Searle's argument for his version of the connection principle is unsuccessful. To put the problem briefly, Searle argues for the connection principle by arguing that no non-mental facts are constitutive of mental facts, and in particular no non-mental facts are constitutive of what Searle calls aspectual shape, the fact that we can think of things under more than one aspect, as we can think of the liquid in a glass as H_20 or as water. Searle then raises the question: what fact about one's unconscious mental states make them have aspectual shape? It cannot be, Searle argues, any non-mental facts about them, facts about their neurophysiology, e.g., because no such facts are constitutive of aspectual shape, which is a necessary condition for intentionality. Therefore, it must be their relation to some other states or properties. The only relation that could do the job is some relation to a conscious mental state, namely, that of being a possible conscious mental state.

The difficulty with the argument is that it equivocates on 'make it the case'. Searle's starting assumption, that no non-mental facts are constitutive of aspectual shape, is a conceptual claim. Thus, in the case of unconscious mental states, if we accept this assumption, we can conclude that no non-mental facts about such states are conceptually sufficient for their having aspectual shape. But this does not mean that they cannot have aspectual shape. For some non-mental facts about such states, e.g., facts about their neurophysiology, so far as anything we have said goes, could be nomically sufficient for their aspectual shape. Or, as far as that goes, they may just have aspectual shape as a brute, unexplained fact. The argument could be repaired by holding that some fact must make it the case that unconscious mental states have aspectual shape other than the fact that they do, and that no non-mental fact is conceptually *or* nomologically sufficient for this. But the second of these claims is clearly question-begging. For it is not advanced as an empirical claim, and if it is a conceptual claim, it is hard to see what could support it other than an implicit appeal to something like the connection principle, which is supposed to be the conclusion of the argument, not one of its premises.

Additionally, even if the argument were successful, it would apply only to states that have aspectual shape. Yet if the connection principle is true at all, one would expect it to apply to all types of mental states. As far as Searle's argument goes, however, although it would be impossible for an individual to have an inaccessible unconscious belief, nothing would bar him from having inaccessible unconscious pains, itches, thirsts, etc. Yet it seems much more difficult to understand how this could be possible than to understand how one could have an inaccessible unconscious belief.

An Alternative Argument

Despite my dissatisfaction both with Searle's formulations of the connection principle and his argument for it, I think the principle is correct in a form stronger than that which Searle gives it, namely:

(CP) Nothing is a mental state unless it is a conscious mental state or it is a disposition to produce a conscious mental state.

The key to seeing why the connection principle in this form is correct is to appreciate the centrality of the first-person point of view in our conception of mental phenomena.

My approach to this will be indirect. We can begin by noting that for a state to be a mental state it must be the mental state of at least one person, and at most one person. The question I want to push is this: what makes a token mental state the mental state of a particular person?[10] We can further divide this question into two parts, one about unconscious mental states and one about conscious mental states.

What makes it the case that a certain conscious mental state is a particular person's mental state? To answer this question we need to specify a relation between a particular conscious mental state and a particular person such that no one else could bear that relation to that mental state. In the case of conscious mental states, the obvious relation is an epistemic one. One has a kind of knowledge of one's own conscious mental states at the time at which they are conscious which no else could have of those mental states. This difference in the kind of knowledge we have of our own and other people's conscious mental states is well illustrated in the methodology of investigations of perception. Contrast the way we find out how a thing looks to ourselves and how it looks to someone else. In our own case, we do not have to ask ourselves for a report of how a thing looks to know how it looks, or to see this by some observation of our behaviour. In the case of other subjects, however, we have no access to how things look to them other than by their reports about it or what differences it makes to their behaviour or performance on various tasks we set them. I will call this kind of knowledge we have of our own mental states which no one else does or could first-person knowledge.[11] One's having first-person knowledge of a particular mental state is sufficient for it to be one's own mental state and sufficient for it to be no one else's mental state. Thus, in the case of conscious mental states, we can say that a

token conscious mental state is *X's* rather than *Y's* because *X* has first-person knowledge of it.

We cannot give this answer in the case of unconscious mental states, because at the time at which they are unconscious, we do not have this kind of knowledge of them. What then makes a token unconscious mental state a particular person's? We can entertain three answers to this question: (1) It is a *sui generis* relation; that a particular token unconscious mental state is a particular person's mental state is a brute fact that admits of no explanation. (2) A token unconscious mental state is a particular person's mental state because it is causally located in his body. (3) A token unconscious mental state is a particular person's mental state because it bears a special relation to that person's conscious mental states. The argument for the third answer will consist in showing that the first two answers are inadequate.

The *sui generis* response can be rejected fairly quickly. If this response were correct, then it would be possible for a token unconscious mental state to bear any combination of causal and epistemic relations to anyone's conscious mental states independently of whose conscious mental state it was. For example, an unconscious mental state which apparently plays a role in your behaviour and produces changes in your mental life, and which is causally located in your body, could still be my unconscious mental state, although it bears no relation to my body, or my conscious mental life at all. This, I think, we will reject out of hand, but since it is a possibility left open by the *sui generis* answer, we must reject that answer as well.

A more plausible answer is that what makes a particular unconscious mental state *X's* is that it is causally located in *X's* body. A state is causally located in *X's* body provided that the intersection of the causal chains in which it is involved is located in *X's* body. The difficulty with this response is twofold.

First, it requires that it not be possible to make sense of *X* sharing his body with anyone else, for if both *X* and *Y* had the same body, then a mental state's being causally located in *X's* body could not ground it as being his, since his body is also *Y's* body. If being causally located in *X's* body were sufficient for it being *X's* mental state, then if *X's* body were *Y's* body as well, it would be *Y's* mental state also. But this violates our starting assumption that every mental state is only one person's mental state.

But it is easy to imagine one body being shared by two or more persons. The relations that hold between a person and his body that makes it his are that changes in it affect his mental states, and in particular his conscious mental states, and that his beliefs and desires explain its behaviour, and more generally, his mental states affect more or less immediately his body. But it seems clear that two different people could bear these relations to one body, either at the same or different times, as is shown by the possibility of conceptualizing cases of multiple personalities as cases of multiple persons occupying a single body. Thus, being causally located in a person's body cannot be sufficient for an unconscious mental state to be his.

The second difficulty is that in attempting to explain what makes a particular unconscious mental state one person's rather than another's by appeal to its being causally located in his body, we have to make sense of what makes a particular body one person's body. The only way to do that is to appeal to its relations to his mental states. But then mental states being causally located in his body could not ground those mental states as his, because what makes it his body is that *his* mental states are causally located in it. This difficulty will afflict any attempt to ground what makes an unconscious mental state a particular person's mental state by its relations to any object the person is contingently related to, for then we will have to specify its relation to the person in terms of its relations to his mental states.[12]

Since the only states a person has essentially are his mental states, no appeal to anything other than his mental states could explain what makes some unconscious mental state his. Since it is obviously circular to appeal to a person's unconscious mental states, and we have rejected the *sui generis* approach, this leaves only the third option, that an unconscious mental state is a particular person's in virtue of its relation to his conscious mental states.

We have so far left unspecified what special relation an unconscious mental state must bear to one's conscious mental states in order for it to be one's own. That it causes a conscious mental state is too weak, since it is possible for one of my unconscious mental states to cause a belief in my psychoanalyst, although that would not make it her unconscious mental state, even if the belief it caused had the same content. The relations between a person's unconscious mental states and conscious mental states must be the sort that it is impossible for anyone else's unconscious mental states to bear to his conscious mental states. No contingent relation could secure this. It must be then that our conception of an unconscious mental state is that it is a disposition of a person to (among other things) have certain conscious mental states.[13] Such dispositions are then individuated in terms of the conscious mental states they manifest themselves as.[14]

Application

Let us now apply this result to the inference theory of perceptual achievement. One of the features of this theory that we noted was that it appealed to unconscious inferences and so to unconscious mental

states. The crucial question is whether these unconscious mental states violate the connection principle. There are at least three reasons to think that they do.

First, note that the condition which the connection principle lays down is very strong. It is not just that the putatively unconscious mental states have a causal effect on my conscious mental states; for this would not distinguish those unconscious mental states that are mine from unconscious mental states that were someone else's. The conception of these unconscious mental states must be of states which are dispositions to produce specific conscious mental sates in the person whose mental states they are. The contents of these states will then be individuated in terms of the contents of the conscious mental states they are dispositions to produce.

The first reason to think that the inferences postulated by the inference theory violate this condition is that our warrant for postulating them is conceived of as independent of the need to verify their occurrence from the first-person point of view. The evidence for their existence is third-person evidence exclusively. This is in contrast to, e.g., the Freudian conception of the unconscious, in which unconscious mental states, though repressed, were in principle manifestable to the consciousness of the intentional agent. This was the ultimate aim of therapy, and was supposed to provide an essential part of the evidence for the theory.[15] The inference theory, however, places no such constraint on the mental states that it postulates. They are thought of as the sort of thing which is independent of the possibility of being manifested in the conscious life of the agent to whom they are attributed. To the extent that they are conceived of in this way, they cannot be mental states at all.

The second reason to think these postulated mental states and processes violate the connection principle is that although they are supposed to play a causal role in the production of conscious mental states, and specifically visual experiences, they are not themselves thought of as dispositions to produce conscious mental states. They are thought of on analogy with a conscious inference of the sort that the theorist might go through in reasoning about how the stimuli impinging on the retina could be used to construct reliable hypotheses about the perceiver's environment. (If my diagnosis below of what has gone wrong is correct, this is not an accident.) But a conscious thought process of this sort is not a disposition of any kind, and would not become so if it were to become, *per impossibile*, unconscious. To borrow Searle's apt metaphor, the inference theory pictures mental states as like fish which can be either at the surface of the ocean or below the surface, so that bringing an unconscious mental state to

consciousness is bringing an item of the very same kind as a conscious mental state into, so to speak, a brighter light. The really deep unconscious mental states, such as those involved in perceptual processing, are, as it were, fish trapped in a underwater cave with no route to the surface. The connection principle in the form I have argued for requires us to reject this picture of the nature of unconscious mental states. All unconscious mental states are dispositions manifestable in part as conscious mental states.

The third reason to think that the inference theory violates the connection principle is that it is intended specifically as a theory of autonomous preconscious mental processing which underlies and explains our conscious perceptions. Its autonomy is required by the fact that its putative product is insensitive to those beliefs about the world and our experience to which we have first-person access, either as presently conscious or occurrent mental states, or as dispositions to have conscious mental states. Thus, that these inferences explain conscious states which are autonomous from occurrent and dispositional attitudes requires that they be autonomous from occurrent and dispositional belief as well, and so violate the connection principle.

If the connection principle as I have formulated it is correct, the inference theory is not just false, but necessarily false, for it violates a necessary condition on anything counting as a mental state for a person, namely, that it is something to which the person has first-person access, either as a conscious mental state or as a disposition to produce a conscious mental state through its manifestations.[16] (We will consider below the strategy of denying that these are to be inferences of the perceiver.)

The first of these points against the inference theory could be met only by admitting that so far we have not the slightest reason to suppose that the inferences exist that are postulated by the inference theory, since we have not the slightest first-person evidence that they do, and would require a radical rethinking of the methodology for verifying such a theory. The second and third points, however, cannot be met without substantially giving up the inference theory.

4. Additional Objections

The argument I have just given aimed to provide an *a priori* refutation of unconscious inference theories of perceptual achievement. But even if one is inclined to doubt that the connection principle is true, there are many other reasons to suspect that the inference theory is on the wrong track. These reasons help to support my claim that the inference theory is deeply mistaken, and also help to support the claim that the

inference theory is in conflict with the connection principle. In this section, I want to have a close look at the kind of explanation the inference theorist proposes to see whether we can make detailed sense of it. I will argue that we cannot.

(1) Do the Inferences Postulated by the Inference Theory have the Right Form *to Be Inferences?*

The first problem which I want to raise has to do with exactly how we are to conceive of these inferential processes. Whether or not one agrees with the connection principle, I think it will be accepted that any inference that is unconscious, even if it is not itself capable of becoming conscious for a given person, is at least the sort of thing which could be conscious. If it is not, then we have no conception of what an unconscious inference is supposed to be. A conscious inference is propositional in form. It involves an agent's beliefs, and is psychologically the acquisition of a new belief on the causal and epistemic basis of another belief or set of beliefs. From the point of view of the agent, this appears as one proposition following from or being supported by another proposition or set of propositions. This description of the form of a conscious inference, together with our requirement that any inference, even if unconscious, be the sort of thing we could imagine being conscious (if not for a given agent, then for some agent), amount to requiring that every inference be representable as propositional in form.[17] When we turn to the sorts of inferences postulated by the inference theorist, however, it appears that they violate this constraint. There are at least two ways this occurs, one at the output end of the inference, and one at the input end.

Let us consider the output end first. The output of such an unconscious inference is literally a visual experience, a way things appear to us. But a way things appear to us is not a propositional representation, and its content is not representable as a proposition. My visual experience of my environment is essentially richer than any beliefs I could have about it. I believe that there is a computer on my desk, that it is rectangular, that is face is grey and smooth, its top white and granular, and so on. But none of this comes close to exhausting the representational content of my visual experience, and this is not due simply to the poverty of my beliefs about the visual scene in front of me. The form of representation itself is different. Since it is a minimal condition on a valid inference that the content of the conclusion be contained in the premises, this means that a visual experience could not be the conclusion of an inference all of whose premises are propositional in form. Since an inference must have a conclusion, and a conclusion must be propositional in form, a visual experience cannot be the upshot of an inference. It is as if you were to offer me as an argument a series of sentences which you call premises, and then hand me a picture as the conclusion. This is just to misunderstand what an argument is. While a picture might suggest a conclusion, it cannot literally be one.

The input end of the inference is if anything even more puzzling, for it is not always clear what we are to suppose the input is like. Sometimes inference theorists talk as if the input were an image on the retina, or a temporal pattern of stimulations of the retina, as, e.g., in the case of the stroboscopic effect. The difficulty with this is that an image on the retina in the intended sense is not a mental image at all, but instead a sequence of irradiations, which is not the sort of thing that could appear as a premise in an argument, any more than a rock or a shadow could. It is no help to shift from the pattern of irradiation to the pattern of firings of rods and cones on the retina. This as well is simply a pattern of physical events, which could exist in the absence of any minds at all. What goes on at the retina may be informational input to the perceptual system in the sense that it is connected in a law-like way with events in our environments, but in this sense of information the warmth of the outside of my coffee cup carries information about the temperature of its contents, and no one would mistake that for a mental state.

To make sense of the idea that an unconscious inference is taking place whenever we perceive our environment veridically, we must make sense of unconscious representations both of general laws connecting what goes on at our sensory surfaces with the nature of our environment, and of particular events. For when we imagine an inference taking place, we are imagining an inference which is valid, and which produces veridical perceptions of our environment, that is, about particular events and objects and processes around us. The physical events at our sensory surfaces, although particular, cannot play the right role because they are not mental. So to make sense of these inferences we must postulate mental representations of these events. There are two ways we can think of these representations. First, we can think of them as perceptual in character. Second, we can think of them as belief-like in character.

If the former, then of course we have the same difficulty as with the conclusion of the putative inference: its content is inappropriate for it to play the role of a premise in an argument. At best we could think of ourselves as forming beliefs on its basis. We have, in addition, if the experience is thought of itself as unconscious, the difficulty of making sense of an unconscious visual experience. The concept of a visual experience seems to be the concept exclusively of a phenomenal

experience, that sort of mental state that, in Nagel's evocative expression, there is something it is like to be in (1979b). Thus, if we take this route with the input, we should treat it as conscious. In any case, if the content of such an experience cannot play the role of a premise in an argument, but at best be a source of information for beliefs whose contents can, we cannot strictly think of it as itself a part of the inference. We must then think of the input to the inference proper as consisting of various beliefs which we have about either what goes on at our sensory surfaces or the content of some perceptual experience.

In the former case, our difficulty is that we lose information in the transition from the perceptual experience to the beliefs which we cannot regain at the putative conclusion of the inference, namely, the final perceptual experience of a scene. This is our first difficulty again. In the latter case, we have this difficulty of course, but also the difficulty of explaining how we come by these beliefs about what is going on at the sensory surfaces. It is evident that at some point explanations in terms of further beliefs must come to an end in a brute fact about the relation between stimulus at the sensory surfaces and our acquiring certain unconscious beliefs. But there seems to be no principle for determining when they should come to an end. In this case, it is at least as reasonable to say that they do not begin at all, and that it is simply a brute fact in the same sense that the stimulation of our sensory surfaces and background conditions produce our conscious perceptual experiences.

If these considerations are correct, then (a) there is no coherent account of the inferences we are supposed to be making unconsciously to arrive at veridical perceptions of the world around us, and (b) there is no reason to postulate them.

(2) *Do Perceivers Have the Concepts Necessary to Perform the Inferences Required by the Inference Theory?*

The second problem has to do with what concepts we would have to attribute to a perceiver in order to think of him as making unconscious inferences of the sort the inference theory postulates. The conceptual resources required by the inferences postulated by the inference theory are those of the inference theorists. But there can be no general guarantee that the conceptual resources of the perceiver match the theorist's. This is obscured by the fact that typically we have in mind the visual experiences of people who have the conceptual resources to understand the inferences that are postulated. But these theories are not supposed to apply just to perceivers sufficiently like the theorists to have such concepts, but also to, e.g., non-linguistic animals which

display a behavioural repertoire sufficient to convince us that they are subject to many of the same illusions and have many of the same perceptual capacities that we do. The evidence is the same in these cases as in the case of human beings. However, it is not plausible to suggest that dogs or fish or pigeons have concepts necessary to entertain the thought that, e.g., the size of an object is proportional to its visual angle multiplied by its distance. Perhaps we have to attribute to dogs and pigeons rudimentary concepts of size and distance to attribute to them visual experiences which represent these, but we have no reason to attribute to them even rudimentary concepts of mathematical operations. Moreover, some of the concepts it would be necessary to attribute to the perceiver to attribute knowledge of principles which must be supposed to be known by him by some inference explanations are not plausibly possessed even by all human beings, for example, the concepts of parallax and luminance, which are invoked in inference explanations of the perception of motion and lightness. For many such concepts the only grounds for attributing them to someone would consist of his speaking competently a language in which some general term expressed the appropriate concept.

This difficulty is a reflection of the fact that the inferences which are postulated to explain conscious visual experiences are treated as autonomous with respect to the perceiver's conscious mental life. This means that no constraints are placed on what concepts can enter into the inferences. But the result of this is that we do not respect the conditions for attributing such concepts to the perceivers. And this shows us that our methodology is mistaken, if we want these inferences to be inferences made *by* the perceiver. For the perceiver could make such inferences only if he was in possession of the concepts which are employed in them.[18]

The Homunculus Response

At this point, it might be replied that it is not the perceiver who makes these inferences, but instead, as is often said, and as I have often said above, it is the perceiver's perceptual system which makes them. It is no accident that explanations of the inference theory so often attribute the inferences to the perceptual system rather than to the perceiver. But this cannot, I think, be a very attractive option once its consequences are appreciated. It commits us to thinking of the perceiver's perceptual system as in effect a different person. For if the perceiver is not making the inference, but the inference is being made, then someone else is making it. If the perceptual system is making the inference, it is a different person from the perceiver. Thus, we would have to

understand the claim of the inference theorist to be that he has not only discovered that unconscious inferences are made in ordinary perception of the environment, but that he has discovered that this inferring is done by a person who inhabits our bodies with us and passes on to us, somehow, his conclusions, as if, in order to know what the world were like, we had to call up someone more closely connected with it to ask him to look outside for us.

This is a version of the homunculus fallacy, the attempt to explain some cognitive function for a agent by postulating a little person inside him who does part of it for him. The first difficulty with this, apart from its *prima facie* implausibility, is that in explaining cognitive capacities it is at best a delaying tactic. This is especially evident if the homunculus is credited with its own perceptual experiences of the world; and it is difficult to see how to avoid this once we have in fact postulated a homunculus, for reasons given below. At some point we must discharge the homunculus if we are to achieve any genuine explanation of perceptual capacities. If we can do so at some point, then it seems there can be no need to postulate a homunculus at any point.[19]

Apart from this, appealing explicitly to a homunculus, while it relieves the pressure on the inference theory from the argument from the unavailability of conceptual resources for the postulated inferences, also removes a good deal of the explanatory point of the theory. Originally, we were to conceive of the theory as giving us an explanation of the perceiver's cognitive abilities. But if the form of our explanation is that the perceiver's perceptions are to be explained by appeal to another person's cognitive abilities, then we have not in fact explained that person's cognitive abilities at all. We have not explained how the perceiver achieves veridical perception by appealing to inferences or knowledge the perceiver has, but to someone else's abilities and knowledge. What we explain now is not how we achieve veridical perception, but how someone else does it for us. This is not to explain our cognitive capacities but to deny that we *have* them.

The homunculus response also undermines the claim that what is going on is that an inference is taking place whose conclusion is in some sense a perceptual experience. For even if we waive the earlier objection that a visual experience cannot be the psychological analog of the conclusion of an inference, we cannot very well allow that an inference could take place in which all of the premises are in one person's mind while the conclusion is in another person's mind. Thus, postulating a homunculus is actually incompatible with the claim that we achieve veridical perception of our environment by an inference from perceptual stimulus to a visual experience, for this requires that we think of the premises of the argument and the conclusion as being

in the same mind, and the homunculus theory denies this. Instead, at best what we have is a half-completed inference in one person's mind and a causal transition from this to a mental state in another's, which is as brute as if it had occurred from the blow of a hammer to the head.

Finally, note that if we treat the perceptual system as a separate person, there is no reason to say that the mental inferences which it is making are unconscious. This is true for homunculus explanations generally. For their not being consciously accessible to us is now no more reason to say that they are unconscious than your mental states not being consciously accessible to me are a reason to call all of your mental states unconscious. While this undermines the criticism of the inference theory based on the connection principle, it also, as we have seen, undermines its explanatory power. That inference theorists want to treat these mental inferences as unconscious shows that they think of them as inferences that the perceiver is making; this is what gives them their explanatory relevance. That they attribute them to the perceptual system reflects their awareness that their autonomy requires that we think of them as not properly mental states of the perceiver at all.

(3) *How Does the Perceptual System Acquire Its Knowledge?*

A third problem is how the system is supposed to know various things about the world which it needs to know in order to make inferences of the sort that result in veridical perceptions. For the inference theory is explaining why we are able to perceive the world veridically in terms of the cognitive powers of our perceptual systems. This is represented explicitly as an intelligent cognitive process which takes information in and produces a perceptual experience as output. The perceptual system does not just make blind inferences, but is in fact attributed knowledge both of what goes on at the sensory surfaces of the individual and in his environment, since it is attributed knowledge of general laws connecting the perceptual stimuli with events, states and objects in the environment. Presumably this is intended in part to explain our knowledge. If the perceptual system did not know what it was about, we could hardly be attributed knowledge on the basis of its products. However, we conceive of our epistemic access to the world around us to be at least partly epistemically mediated by our visual experiences, and we have no conception of how else we could gain knowledge of the world around us. Thus, if we attribute knowledge of the world to our perceptual system, then it must either be derived from knowledge that we have independently, or the perceptual system itself must possess the kind of epistemic access to the environment that we do – that is, it must independently have perceptual experiences, etc.

Neither option is acceptable. If we take the first option, there are two problems. First, if the knowledge we have is to depend on knowledge the perceptual system has independently, then this appeal deprives both us and our perceptual system of any knowledge of the world, since we would presumably have such knowledge only if our perceptual system did. Second, most of us do not have the sort of knowledge that is attributed to the perceptual system at all. The rules which our perceptual systems are supposedly employing in inferring what our environment must be like are supposed to be uncovered not by first-person reflection on what we already know, but by third-person investigation of how our perceptual system achieves veridical perception of the world around us. If we had first-person access to such rules, then there would be no need for psychologists to undertake to discover what they were. If we take the second option, then it is clear that we are attributing to the perceptual system its own perceptual system, and the regress we noticed above is in full swing.

(4) Can There Be Any Evidence for the Inference Theory?

The last objection I want to raise in this section is a methodological one. Philosophers are often interested in claims about what is possible or not even when there is no empirical method for discovering whether or not the hypothesis in question is true. One of the things that sets the sciences apart from philosophy is that scientists are not interested in hypotheses which are not empirically confirmable or which are not needed to account for our observations. Methodologically, scientists are verificationists (thus, the relative lack of interest among physicists about the question of hidden variables in quantum mechanics except as that can be shown to make an experimental difference). In our discussion of the inference theory above, it appeared that it met this criterion for being a scientific theory because we were able to show that in some cases it was superior to the stimulus theory. This appearance, though, is illusory because we know that in principle there is no need to introduce the notion of unconscious inferences to explain our perceptual capacities. All that is necessary is that the actual process that produces our perceptual experiences connect in a law-like way features of our environment with the representational content of our experiences. The process which does this can be entirely physical up to the point at which we know that some mental state occurs. We know that from our first-person access to those mental states. Not only would a physical story be adequate, we are committed to there being such a story, for even if we were to accept that there were unconscious inferences occurring in the production of conscious

visual experiences, we suppose that this would not be an irreducible fact about the process, but would itself be instantiated or realized in an underlying process that had a purely physical description. So we are committed to saying that there is an adequate account of how we get from sensory stimuli physically described to our perceptual experiences which does not postulate any unconscious inferences. Thus, the postulation of such unconscious inferences is gratuitous. This shows that, given our commitment to a physical basis for all thought, we could not in principle have a reason to postulate inaccessible unconscious inferences on the basis of the evidence that we achieve veridical perception of the world around us, for we know (or are committed to holding) that this admits of an explanation that does not require inaccessible unconscious inferences. The reason the same argument does not apply to conscious states, or dispositions to produce conscious states, is that these are not theoretical entities for us, but epistemically primary. We have independent reason to believe that they exist. But the only reason there could be to believe that inaccessible unconscious mental processes existed would be that they were explanatorily indispensable. We are committed to their being explanatorily dispensable, so it follows that we can have no reason to think that they exist; hence, they can be of no interest in an empirical theory. Such unconscious inferences, even if we could make sense of them, would have no place in a scientific psychology.

6. Diagnosis

If the arguments I have given are correct, then the inference theory is deeply misconceived. How then does it come to seem so compelling?

It is instructive to begin by considering again the comparison of the stimulus theory and the inference theory. In some of the cases we considered, there seemed to be a clear choice between the inference theory and the stimulus theory. If the mechanism were as the stimulus theory described it, then it would be inappropriate to describe the process as involving unconscious inferences. But this should be puzzling. For *prima facie* we ought to be able to treat the mechanism postulated by the stimulus theory as a matter of unconscious inference as well. For example, consider the stimulus theory of size constancy which holds that size constancy is a matter of the ratios of the sizes of images on the retina. Why should we not represent this way of achieving size constancy as a matter of a very simple inference from the fact that the ratio of the size of one image to that of another at one time is the same as it is at a later time, and a general rule that holds that objects do not change in size as long as the ratios of their visual angles remain constant? Again,

consider the case of the waterfall illusion. Why should we not say that the perceptual system infers from the contrast in the pattern of cell discharges that the motion of the scene has not reversed its direction? I think our reactions to these cases contain important clues to the appeal of the inference theory. In the first case, part of what is important is that the putative inference involved takes into account only one factor, and, hence, seems more mechanical than in the case of the process postulated by the inference theory, in which distance and visual angle are independently determined and then brought under a law to determine size. In the second case, this factor is at work, but there is an additional important factor as well. That is that the inference which might be postulated to explain the waterfall illusion is not plausibly thought to be the result of the application of an assumption which would ordinarily be correct and result in veridical perception. Thus, we postulate an inference to explain a perceptual effect only when it is conceived to rest on a true *ceteris paribus* law connecting some feature of the perceptual stimuli with some feature of the environment.

This is a direct consequence of the question that defines our inquiry. We want an account of the process that produces *veridical* perceptions. A necessary condition for this is that the process we describe connect up representations with what they represent in a law-like way. Thus, our goal is to identify a set of laws L_1 that correlates properties of the environment with properties of the stimulus patterns at the retina, and a set of laws L_2 that correlates properties at the retina with representational properties of our visual experiences, so that L_1 and L_2 jointly entail a set of laws L_3 that correlates features of our environment with representations of them. Thus, an inference theory which represents the laws in L_2 as being instantiated by a process of reasoning must attribute to the perceiver or perceptual system assumptions about true laws of the sort that would go in L_1. That these laws are, of course, *ceteris paribus* laws allows for the possibility of perceptual errors and illusions. Thus, the difficulty with the postulated inference in the case of the waterfall illusion is that the assumption that would be attributed to the perceptual system is not a true *ceteris paribus* law connecting features of the perceptual stimuli with features of the environment, while the kind of mechanism we are interested in instantiates a true law of that kind. The trouble with the stimulus theory of size constancy is of a different kind. Here it is simply that there does not seem to be much of a mechanism required and so not much need to think of the mechanism as instantiated in inferences that the perceptual system makes, for lack of any other hypothesis.

But in neither of these cases do we have a reason to deny that the process that the perceptual system is undergoing involves an unconscious inference. The reason to deny this would be, in the case of the waterfall illusions, that the perceptual system makes only correct assumptions about the laws that connect the environment with stimuli at the sensory surfaces; in the case of the ratio theory, it would be that the perceptual system only makes relatively complicated inferences. The evidence we have, however, provides no reason to make these assumptions. As far as the evidence we have goes, we could attribute to the perceptual system a series of false assumptions which fortuitously result in correct conclusions, or very simple inferences from simple features of the perceptual stimuli.

What then explains the pattern of inferences actually attributed to perceivers? What has happened, I think, is the following. We have started out with the assumption that we perceive our environment by and large veridically, and so we want an account of the perceptual system that connects sensory stimuli with veridical representation. We give a set of laws L_1 that connect features of our environment with sensory stimulation. We hypothesize that the laws L_2 that the perceptual system instantiates take account of the sensory stimulations cited in the previous laws. Where it seems that we need to give an account of the *mechanism* which instantiates these laws, given that we want to match representations with what they are of, and think of this as the goal of the perceptual system, *we adopt the expedient of assigning knowledge of the laws in L_1 to the perceptual system, and treat the mechanism as a matter of a rational agent inferring from the laws in L_1 and the stimuli to the nature of the environment.* The mechanism that instantiates the laws in L_2 then becomes a psychological mechanism.

It is clear that this is invalid. Why is it so attractive? There are a number of reasons. First, it is a mechanism that we understand and whose postulation does not require us to do any detailed investigation of the neural mechanisms underlying the process. Second, since we think of the perceptual system as having as its function the production of veridical experiences, it is easy for us to overlook the distinction between the theorist's point of view and the point of view of the perceptual system. That is, thinking of the perceptual system as having a function suggests that it has its own goal, hence, its own point of view; once we have got that far, it is easy to identify its point of view with our own. We can infer, given the laws that connect the environment with surface stimuli, what experiences are needed to have veridical perceptions of the environment. The perceptual system then, we think, having this as its goal, must do the same. This bit of transference is clearly a mistake. What lends to the confusion is that while in the case of many biological functions the output is not a conscious experience – think of the digestive system – in the case of the visual

system it is. When we then think of the perceptual system as having veridical perception as its biological function, since the goal itself is an intentional state, it is easy to treat this as an intrinsic goal of the system, and natural to think of the goal driven process by which it is produced as mental as well.

That this is a mistake is brought out in the following two thought-experiments. In the first, suppose that we have added onto the perceptual system a mechanical device that takes into account a number of different factors in pre-processing the image that is projected onto the retina. For example, suppose that it filters out images on the basis both of size and the wavelength of light they reflect. The process that now results in veridical perception includes some additional stages of processing: would we ascribe additional unconscious inferences to the self or the perceptual system? In this case the answer seems clearly to be "No." But our position with respect to the perceptual system is no different from our position with respect to an apparatus designed to pre-process information before it arrives at the retina. In the second, we can suppose that we have designed a machine that can keep track of the movement of objects through keeping track (entirely mechanically) of its own movement, and the movement they generate by changes on a two-dimensional array of light detectors. The machine keeps track of this in the following sense: it has a mechanical model within it of the space around it in which various elements in it correspond to elements in the environment, and their movements to movements of the objects in the environment. We do not, I think, feel any inclination to say that it is making inferences. Suppose we now add a Cartesian soul which has visual experiences of such objects moving as a causal result of this mechanical process. It is obvious that this changes nothing about the process itself.

Similarly, in the case of some physical processes causing in one an experience which is not representational in character, we feel no inclination to say that the process that produced it was a cognitive process, even though the process may be functionally quite complex, as in the case of the digestive system. It could even be at some level of description functionally identical to the perceptual system. But when an upset stomach causes stomach ache we do not suppose that the process that produces the unwanted experience is itself mental. It is only when we start out with a system that has mental states and a process that produces representational states that we feel inclined to suppose that not only the product of the process but the process itself is mental in character.

There are five further possible sources of the seductiveness of the inference theory. The first is the tendency to think of perceptual illusions as cognitive failures. If illusions are cognitive failures, then we are forced to think of perceptual illusion as a matter of some failure in taking proper account of the world. But then this must be represented as our having made some false assumption, and the illusion is treated as an instance of mistaken reasoning. The mistake here is to treat all errors as if they were errors of reasoning.

This is connected with a second possible source of confusion, which is to take the verb 'to perceive' to be analogous to an action verb, and so to be a cognitive achievement verb. If perceiving is something we *do*, in the same sense in which we reason or act, then it must be the result of a cognitive process just as reasoning and acting are. Thus, a proper explanation of perception must bring in a cognitive component. The mistake here is to treat perceiving as if it were something we did in the sense in which we write a paper or solve a puzzle. The sense in which to perceive something is to do something is the same sense in which to breath, or to perspire, or to dream is to do something. None of these implies that what is done is a result of a cognitive process. Perceiving is more like feeling pain that it is like drawing a conclusion.

The move to treating the process by which veridical perception is achieved as an inference is probably also aided by a failure to distinguish between two senses in which we talk of information-processing. There is first of all causal information-processing. In this sense, a system processes information, provided that its states are connected in a law-like way so that a potential observer who had knowledge of the laws governing the transformations of the states could recover information about what happened earlier from later portions of the process. In the second sense, information-processing is just what we consciously do when we reason from our beliefs; in this case 'information' means 'representation'. In the first sense of information-processing, the visual system processes information, but so equally does the digestive system, the immune system and the solar system. In the second sense of information-processing, no one would suppose that the digestive system processes information, but when we turn to perception, since the end of the process is a perceptual state which carries information in the *second* sense, we can be led to suppose that the causal information process must itself carry information in the same sense.

A fourth possible source of confusion is failure to observe the distinction between two senses of rule-following, descriptive and normative. A system follows a rule descriptively, provided that its temporal development is correctly described by the rule. A system follows a rule normatively, provided that its temporal development is explained by its knowing a rule and acting in accordance with the rule out of the intention to follow it. The perceptual system follows

descriptively the rule: perceive one thing to be larger than another if its distance is greater but it has the same visual angle. When we think of the output as a perceptual state, and the system as having veridical perception as its biological function, it is easy to read the rule as normative rather than simply descriptive.

A final possible source of confusion is the tendency to conflate two different notions of intelligent behaviour. In the first sense, we speak of a system acting intelligently in the sense that it acts *as if* it were intelligent. This can be a context sensitive classification. Thus, we contrast, e.g., a thermostat which turns on the heater when the temperature drops below 68 degrees, with an "intelligent" thermostat which takes into account, say, the humidity in addition to the temperature. In the second sense, we mean a system is genuinely intelligent in the sense that it has intentional states and engages in explicit reasoning and deliberation to produce behaviour appropriate to its goals. The perceptual system is certainly intelligent in the first sense; that is why the ratio theory of size constancy is inadequate. Since the output of the perceptual process is a perceptual experience, it can be easy to move from thinking of the perceptual system as behaviourally intelligent to thinking of it as genuinely intelligent.

7. Reconstruction

What can be salvaged of the inference theory? Here I think two things can be said. (1) First, if I am right in my diagnosis of the inference theory, there is a stage in the investigation in which we make genuine empirical hypotheses about how the perceptual system achieves veridical perception of the environment. This is in part a matter of discovering exactly which of the features of the perceptual stimuli are both correlated reliably with features of the environment, on the one hand, and with representations of those features, on the other. It is this which the empirical evidence has a bearing on. This is essentially a question from the design standpoint about which features of the sensory stimuli are the ones the perceptual system is causally attuned to in the production of visual experiences. This is an important and interesting empirical question, and conclusions about this can be extracted from the debate between the inference theory and its rivals.

Second, the inference theory is not just about what features at the sensory surfaces are relevant to the production of a veridical perceptual experience but also about the process by which those features produce that experience. Although I have argued that those processes cannot be mental processes, this does not mean that there is nothing of value in the hypotheses of the inference theory, for they can be

treated as hypotheses about the functional specification of the process that produces the visual experience. For example, the inference theory holds that we achieve size constancy in visual experience by the unconscious assumption that Emmert's law is true and by the belief that a certain object is a certain distance and that it subtends a certain angle on the retina. We can convert this into the hypothesis that size constancy in visual perception is a function, given by Emmert's law, of two things: (1) the visual angle an object subtends and (2) the features of the stimuli that are responsible for those states in the perceptual system that indicate the distance of the object. Thus, we give a partial functional characterization of the physical process whose upshot is that a certain object looks to be a certain size. What we aim for, on this view, is a functional characterization of the perceptual system which treats events at the sensory surfaces as input and veridical representations as output. The fact that the output is veridical representation requires that the selection of the input and the functional organization of the perceptual system be such that in our actual environment our perceptual representations are by and large correct. Such a functional specification selects from among a number of possible systems, and provides a guide to an investigation of its neurophysiological realization. What it does not do is provide us with a distinctively psychological explanation of how veridical perception is possible. It provides a psychological explanation only insofar as the output is psychological, or the causal processing itself is influenced by conscious psychological processes. The same kind of story can be told about the immune system; this does not make the account of the functional organization of the immune system a psychological account.

Thus, I advocate a retreat from the explicitly intentional language of the inference theory to the functional structure any such inferences must assume.

8. Objections to My Argument

Objection 1. Psychologists are just extending in a principled way the meanings of the terms 'mental' and 'inference', so any objection based on the concept ordinarily expressed by these words does not apply to the inference theory.

Reply. It is not clear that this is the intent of any actual inference theorist, but in any case reformulating the inference theory explicitly as offering an extension to the ordinary conception of the mental would achieve at best a Pyrrhic victory. I can extend the meaning of any word by adding a disjunct to its definition, but then when I argue on the

basis of the extension that some item falls under the concept now expressed by that word I do not thereby show that we have discovered that something falls under the concept previously expressed. If I say that henceforth I will call anything that is a human being or a chair "a human being," and then assert, "I am sitting on a human being," I have not discovered a new and startling fact but relabelled an old and mundane one. This discovery would not call, for example, for new legislation to protect the rights of a previously unrecognized segment of humanity. Similarly, to suggest that the inference theorist is or ought to be just extending the meaning of 'inference' is to suggest that the theory is of very much less interest than it seemed to be at first. It might still be urged that the extension is a principled one, and that therefore the theory is making a substantive claim. But if the argument I have given is correct, the extension would consist of including in the extension of 'inference' information-processing in the sense in which a sunflower or a prism processes information; in this sense virtually every causal process is a mental inference, which just trivializes the claim that the perceptual system is making unconscious inferences. There is, furthermore, a point to not changing usage as this strategy urges: in doing so we obscure distinctions important to a clear understanding of what's special about having a mind.

Objection 2. Sometimes conscious knowledge or belief influences the character of perceptual experience. For example, in the perception of ambiguous figures such as the Necker cube, a suggestion to a subject about what the figure is before he views it will often determine how he sees it. Since the input here is clearly mental, the process which selects which one of the two figures the figure is naturally seen as must be mental as well. A similar phenomenon is the autokinetic effect. When a subject views a stationary point of light in a dark room, often it will be seen as moving; and whether it is seen as moving and in what direction is susceptible to suggestion. Thus, it appears that a conscious belief plays a role in how a scene is perceived; this requires that the process which produces the perception be a mental process.

Reply. That what we consciously think we are seeing or likely to see should have a causal effect on the process that generates our visual experience does not require that the process itself be a mental process. If what I think causes my ulcer to act up, which causes me discomfort, it does not follow that the causal process that led to that is a mental process. Similarly, when I type a sentence at my keyboard, I do this because of my intention to write a particular sentence; but it does not follow that each movement of my fingers is itself intended by me to

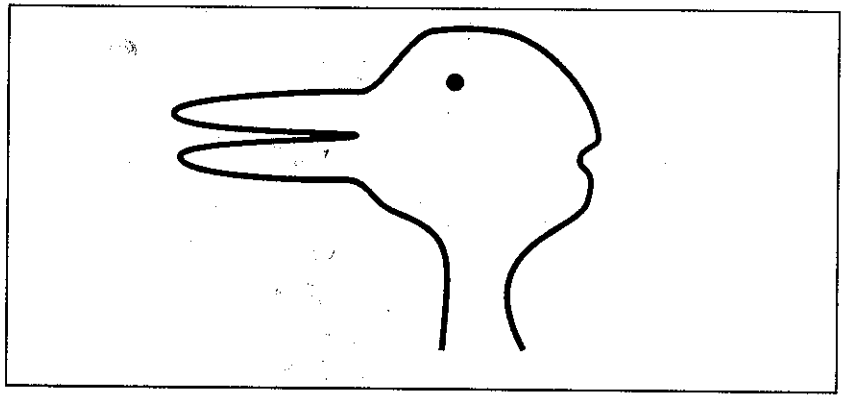

Figure 2.6: The Duck-Rabbit

move as it does or that each movement is the result of an unconscious process of the same sort as the conscious process which results in my typing the whole sentence. What makes us want to say that in this case there is a kind of mental process is that the outcome of that process is one of two possible representations of an ambiguous figure, and that we have a model for what it is like to disambiguate some figures or sentences by a conscious process. But the model is inappropriately applied in this case, because it is not possible to bring to our awareness any such mental process. Furthermore, clearly this kind of case should be expected even if one thinks that the process is entirely non-mental. For given that it is advantageous to see objects as falling under certain kinds on the basis of a few cues about them, it will be possible to produce ambiguous figures, and since it is advantageous to be primed to see what one expects in one's environment on the basis of as few cues as possible, one would expect that the experiential upshot of seeing an ambiguous figure would be causally sensitive to one's expectations about what one is likely to see.

Objection 3. Many of our perceptual abilities obviously depend upon learning. This is clear, e.g., in the case of many pictorial cues to relative depth, such as interposition patterns. Learning is also clear in cases of perceptual adaptation, as in the case of subjects adapting to lenses which distort the image projected onto the retina. A simpler case in which learning is clearly important involves ambiguous figures such as the familiar duck-rabbit figure (Figure 2.6), which would clearly not present itself to us as either a duck or a rabbit unless we were familiar with and could recognize features of both ducks and rabbits. Someone unacquainted with animals altogether would not see such a drawing as representing anything. But if such perceptions depend

upon learning, then it must be that at some time one came to recognize an object by an explicit process of detecting features and consciously bringing an object under a concept as a result. Similarly in the other cases. When one becomes skilled at this, the process takes place below the level of consciousness. But this is no reason to deny that it is a process of the same type.

Reply. First, that a skill or ability to recognize something is acquired does not show that the acquisition of it involved an explicit process of taking into account features of an object or scene and recognizing that the object or scene falls under a concept in virtue of having those features. But even if the process by which the skill is acquired did involve initially taking account explicitly of features of a scene in deciding that it fell under a concept, that does not require that we think of the skill once it is acquired as a matter of an explicitly represented inference or recipe that is followed that has dropped below the level of consciousness. Representations and skills are in different logical categories. This can be seen by noticing that in the exercise of any complex skill that involves explicitly represented parts, one must have the ability to carry out the parts. If each part in turn had to be broken down into a series of explicitly represented instructions, then since each instruction could be carried out only if one had the ability to do it, we would have an infinite regress. The parallel with acquiring a skill such as playing tennis is useful here. One can receive instruction in how to serve in tennis. But when one has acquired a good serve, one does not any longer go through those instructions, either consciously or unconsciously. The acquisition of a skill is the process of dropping the need for any explicit representation of the task below the level of representing the goal toward which the skill is applied. In the same way we can acquire skills in recognizing objects without supposing that there is any processing or inference which goes on below the level of consciousness, even if originally we learned to recognize an item with the help of explicit instruction.

Objection 4. If recognitional abilities are acquired, then this implies that past experience is taken into account in our present perceptual processing. This means that some physical trace of the past experience is left in the brain, which is a memory trace, and that when we see something what we see it as is a result of a very rapid unconscious search through these memory traces. This search can be conceived of as either a parallel search or a serial search. Moreover, it is clear that there is empirical evidence which bears on this question, so that it cannot be ruled out *a priori*. For example, in one experiment, subjects were

given a list of items to memorize. Thereafter, they were shown an item and asked to say whether it was on the list. The amount of time it took subjects to answer was proportional to the number of items on the original list they were asked to memorize, thus bearing out the hypothesis that a very rapid but serial search through memory was being performed. This process is experimentally confirmed, and obviously too fast to be accessible to consciousness. Therefore there cannot be any in principle objection to unconscious inferences in explaining cognitive abilities.

Reply. All that is implied by the fact that recognitional abilities are acquired as a result of our experience is that they are causally sensitive to past experience. This is a genetic fact about them, not a fact about how the recognitional ability is realized. The mistake in the above argument occurs in the movement from 'physical trace' to 'memory' in the sense of a representational state. In the sense in which a physical trace is a memory, a tree remembers which years were wet and which were dry, because this leaves a physical trace in the rings in the tree's trunk. But it does not follow that the tree represents certain years as wet and other years as dry. The experiment designed to distinguish between serial and parallel processing in recognitional abilities requires careful treatment. This is not a good test case for the inference theory because in the experiment described the representation of the different items on the list are straightforwardly accessible to consciousness in the sense that once the list is memorized, the subject has a disposition to form a conscious thought that an item was on the list as a result. Thus these representations are not inaccessible unconscious representations of the sort that might be involved in basic perceptual recognitional abilities. It might be thought, nonetheless, that there is a process of unconscious reasoning going on, an unconscious comparison of the item the subject is shown with an unconscious mental image or representation of each successive item on the list. This must be rejected on my account because such a process is conceived of as involving states which are not simply dispositions to produce conscious mental states but just like a conscious process of inference except for being unconscious. The reply to this is that there is certainly no requirement that the experimental results be understood this way. They are perfectly compatible with the claim that there is no unconscious process of inference at all; the results tell us something about how our recognitional abilities are implemented, but do not require us to think of them as implemented in a process of unconscious inference. What the experiment tells us is something about how our dispositions to have certain conscious thoughts and our desire to perform a

certain recognitional task interact to produce a conscious judgement about whether an item is or is not on a memorized list.

Objection 5. Toward any visual experience we can take a stance which we can call proximal perception.[20] This is the standpoint an artist takes toward his experience when he wants to paint so that the painting will look like the scene. Thus, in the case of the Ponzo illusion, it is possible to see or experience the scene, even if with some difficulty, simply as two lines on the page which are converging at the top of the page, and the two circles as at least more nearly the same in size. This shows that the input to the process that produces the appearance that the upper circle is larger is accessible to consciousness, and so that the input is actually a mental state. If the input is a mental state and the output is a mental state, then the process that takes us from the one to the other must be a mental process.

Reply. There are two mistakes in this objection. The first is that we are not licensed to infer from the fact that how things appear to us can change if we adopt the stance of proximal perception that things appear to us that way in any sense when we do not. So we are not licensed in drawing the inference that we see things that way all the time unconsciously and then infer unconsciously from the character of that experience to how things should be represented. On the other hand, if we are aware in some sense even when perceiving a road receding into the distance that the edges of the road appear to be converging toward the top of the visual field, then our perception of depth is not a matter of an unconscious inference from unconscious input to a conscious experience. In this case it might still be said that an unconscious inference takes place. But this is the second mistake. The fact that we make a transition from one mental state to another mental state does not require that the transition be mediated by a mental process. If it did, then since a mental process is itself a sequence of mental states, we would immediately have a vicious regress.

Objection 6. If you deny that the inference theory is coherent, then you will have to deny as well that we engage in unconscious mental processing in cases such as disambiguating sentences in natural languages. For example, consider the following sentence:

Bud and Pearl saw the Great Lakes while they were driving to Canada.[21]

We understand the antecedent of 'they' in this sentence to be Bud and Pearl rather than the Great Lakes without any conscious awareness of disambiguating it. But it is overwhelmingly plausible that we do this

on the basis of our knowledge that interpreting 'they' as referring to the Great Lakes would require an interpretation someone could intend only if he were seriously confused about the nature of large bodies of water. But if your criticism of the inference theory is correct, you are committed to saying that we do not arrive at this interpretation by any mental process whatsoever.

Reply. That we disambiguate this sentence on the basis of our knowledge about the relative plausibility of someone who uttered it intending 'they' to refer to Bud and Pearl rather than to the Great Lakes is not incompatible with my account. This knowledge is clearly not ruled out by the connection principle because if one points out to someone that 'they' has two interpretations, one can explain with little difficulty why it is implausible that someone would utter such a sentence intending 'they' to refer to the Great Lakes rather than Bud and Pearl. What my account denies, however, is that when we disambiguate such expressions on the basis of our knowledge, what we do is to go through a very rapid unconscious inference in which we entertain each of the interpretations and reject one on the basis of its implausibility. Rather, our dispositional knowledge causally conditions how we understand the sentence. But this does not require that there be an extremely rapid unconscious inference.

Objection 7. The sorts of inference theories you have considered here are unsophisticated. More sophisticated theories (which provide detailed accounts of how the incoming signal is processed stage by stage to generate a 3-D image from the 2-D array of input at the retina) are not subject to these objections.

Reply. This objection mistakes my criticisms for empirical criticisms of the inference theories I have considered which claim that these inferences are too coarse-grained to provide an adequate account of how our perceptual mechanism works. The objections that I have raised do not depend upon how detailed or sophisticated the inferences postulated are but only on their being genuine mental processes just like conscious processes except for being independent of our conscious mental lives.

Notes

1 A prominent recent defender of the theory of unconscious inferences in perceptual processing among psychologists is Irvin Rock (1983, 1984), whom I will use for purposes of illustration. Among Philosophers, Jerry Fodor is well-known for defending this position (1983).

2 I make no pretence in what follows to provide an exhaustive account of the kinds of experimental evidence bearing on the perception of size and motion that psychologists have accumulated and strategies for explaining it in terms of unconscious inferences they have employed. That would be an enormous undertaking. My aim is to provide enough of a sketch of the kinds of evidence available and the kinds of explanation which have been offered to fix our subject matter and give substance to the criticisms and diagnosis that I will offer.

3 For example, the central criticism I make will apply to all computational theories of cognitive capacities which (a) treat computations as operations over genuine representations and (b) treat them as independent of our conscious mental lives.

4 In this sketch of an inference account of visual perception of size constancy, we have taken for granted that the visual system can determine the distance of an object from the observer. The explanation of how it does this can be expected to take into account many different facts, such as recognition of familiar objects whose sizes we know, information about the convergence of both eyes on an object to keep it in focus, the tendency of the eye to accommodate for the distance of an image to keep it in sharp focus, retinal disparity, motion parallax and pictorial information. The explanation of how the visual system determines distance can be given in terms of unconscious inferences from perceptual cues about the relative distances of objects in the environment. For our purposes, nothing essential is left out in concentrating on the stage of processing which moves from distance information and visual angle to relative size.

5 It is striking that this account of how the inference theory would apply to the perception of motion can be given just from a description of the perceptual problem. This should make us suspect that in some cases the empirical content of the inference theory does not exceed the description of the problem itself.

6 As Irvin Rock puts it at one point, "According to this theory, apparent movement is a solution to the problem posed when object *A* disappears in one place in the scene and another object, *B*, suddenly appears in another place. After all, this sequence is quite similar to real motion, particularly when it is rapid" (1984, 195). And "Thus, in an apparent-movement display, when the conditions mimic those of real, rapid motion, entailing sudden disappearance of an object in one place and its reappearance in another, our perceptual system makes the plausible inference that the object has moved." (1984, 196)

7 The point of this paragraph is not to endorse this alternative explanation, which may be disputable, but to show that empirical findings can apparently bear on the correctness of the inference theory. For this purpose it is not necessary to show that theory is false, but only that it could be.

8 The connection principle is not new; most classical philosophers apparently held it in one form or another. For example, Descartes apparently held that no state is a mental state unless it is a conscious mental state. Thus: "As to the fact that there can be nothing in the mind, in so far as it is a thinking thing, of which it is not aware, this seems to me to be self-evident. For there is nothing that we can understand to be in the mind, regarded in this way, that is not a thought or dependent on a thought. If it were not a thought or dependent on a thought it would not belong to the mind *qua* thinking thing; and we cannot have any thought of which we are not aware at the very moment when it is in us" (1647/1985, 171-72). Thomas Nagel is a contemporary adherent along with Searle: "Not all mental states are conscious, but all of them are capable of producing states that are" (1979c, 188).

9 The first formulation, that every mental state is or is potentially a conscious mental state, suggests a picture of the relation between conscious and unconscious mental states that makes it difficult to see how the connection principle could be a conceptual truth, for it suggests that unconscious mental states are just like conscious mental states except for being unconscious. In a metaphor that Searle uses himself, this is the conception of unconscious states as like fish below the surface of the sea, which have only to be brought to the surface. On this conception, it is puzzling why we must think of an unconscious mental state as essentially tied to a conscious mental state. Perhaps we must think of it as the sort of thing which could become conscious, that is, as falling under a mental type such as belief, which a conscious mental state could also fall under, but this leaves it open to think of it as the sort of thing which for a particular individual is completely inaccessible to him. The second formulation can be read in either of two ways. On the first reading it is simply the claim that the type under which an unconscious mental state falls must also be a type under which a conscious mental state can fall. On the second reading, it is the claim that necessarily every unconscious mental state of a given individual is a potential conscious mental state for that individual. Both of these readings are distinct from the last form because neither says anything about "the ontology of the unconscious" and neither entails that what it is for a state to be an unconscious mental state is for it to be capable of causing a conscious mental state. The last formulation is too weak, for reasons which I give below.

10 The question is really more general than this, if there can be creatures with mental states which are not persons. For convenience, however, I will continue to speak of persons.

11 To say that we have a special kind of knowledge of our own conscious mental states is not, however, to claim that we are incorrigible or infallible about them. It is therefore not an objection to this distinction that no certainty attaches to our pronouncements or beliefs about even our conscious mental states.

12 It might be objected at this point that we can first identify a body as a particular person's body by reference to its effects on his conscious mental life, and then identify his unconscious mental states in terms of their being causally located in his body. But this does not avoid any problems, since if a body is identified initially by its contingent relations to a person's conscious mental states, then there can be no in principle bar to more than one person bearing those relations to that body.

13 We must think of it also as a disposition to produce behaviour in conjunction with one's other beliefs and desires.

14 The dispositions which we identify as unconscious mental states are dispositions of the person whose states they are. Thus, although it may be true that the wall is in a certain sense disposed to cause a visual experience in me, that disposition is not my unconscious mental state because it not a disposition of me. The wall could cease to exist without the visual experience or the perceiver ceasing to exist.

Note that the conception of unconscious mental states presupposed in much of the argument, as states just like conscious mental states only unconscious, has now been abandoned in favour of conceiving of unconscious mental states as being, as it were, mental by courtesy of their power to produce conscious mental states.

15 The Freudian theory, however, as traditionally conceived, violates the connection principle for other reasons, for it conceives of unconscious process as just like conscious processes, rather than just as dispositions to produce conscious mental states.

16 Joseph Tolliver has suggested helpfully that one could think of this as a secondary qualities theory of unconscious mental states.

17 This is in fact the way that inference theorists represent the inferences that they attribute to a perceiver or his perceptual system. Here is an example given by Rock (1983, 274): "*Major premise:* An object's visual angel is inversely proportional to distance. *Minor premise:* Visual angle is 1 degree (producing a particular perceived extensity); distance is 50 feet (producing a particular perceived distance). *Conclusion:* Object is equivalent to one that would yield a visual angle of 25 degrees at 2 feet (or 5 degrees at 10 feet, etc.)." As I note in the text, this cannot be correct, because strictly speaking the "conclusion" is supposed to be a visual experience.

18 Another aspect of this disconnection between the usual procedures for attributing concepts to someone and the procedures of the inference theory is that the attribution of beliefs and assumptions in the inference theory is radically indeterminate. There are clearly many different sets of premises which could be attributed that would have the same conclusions. But since the attribution of such inferences is not constrained by any evidence that the perceiver has the concepts needed to entertain them, or by the perceiver's awareness of such inferences, any inference which produces the right result compatibly with the constraints on the input and the constraint that the perceptual experience be such that it is by and large veridical in the perceiver's environment is as good an any other. There is no objective ground to choose between them.

19 One might attempt to defuse this objection by arguing that as we go down levels we get to dumber and dumber homunculi so that we really are achieving some explanatory progress. However, in this particular case, we have a homunculi that has conceptual resources which may be more sophisticated in some ways than those of the perceiver, and there is no reason to think as we go down levels the lower levels will involve any less sophisticated concepts.

20 I borrow the term from Rock (1983, 1984).

21 Adapted from an example by George Miller.

References

Anstis, S.M., and R.L. Gregory (1964). The after-effect of seen motion: The role of retinal stimulation and eye movements. *Quarterly Journal of Experimental Psychology*, 17, 173-74

Descartes, R. [1647] (1985). *Fourth Replies*. Translated by John Cottingham, Robert Stoothoff and Dugland Murdoch. In *The Philosophical Writings of Descartes*. Cambridge: Cambridge University Press

Fodor, J. (1983). *The Modularity of Mind*. Cambridge: MIT Press

Gibson, J.J. (1950). *The Perception of the Visual World*. Boston: Houghton Mifflin

Helmholtz, H. von. [1867] (1962). *Treatise on Physiological Optics*, Vol. III. Edited and Translated from the 3rd German edition by J.P.C. Southall. New York: Dover Publications

Ludwig, K. (1993). A dilemma for Searle's argument for the connection principle. *Behavioral and Brain Sciences* 16: 194-5

Nagel, T. (1979a). *Moral Questions*. Cambridge: Cambridge University Press

———. (1979b). What is it like to be a bat? In T. Nagel (1979a), *Moral Questions*. Cambridge: Cambridge University Press

———. (1979c). Panpsychism. In T. Nagel (1979a), *Moral Questions*. Cambridge: Cambridge University Press

Penrose, L.S., and R. Penrose (1958). Impossible objects: A special type of visual illusion. *British Journal of Psychology* 49: 31-33

Rock, I. (1983). *The Logic of Perception*. Cambridge: MIT Press

———. (1984). *Perception*. New York: Scientific American Library

Rock, I., and S. Ebenholtz (1959). The relational determination of perceived size. *Psychological Review* 66: 387-401

Rock, I., and S. Ebenholtz (1962). Stroboscopic movement based on change of phenomenal rather than retinal location. *American Journal of Psychology* 75: 193-207

Searle, John R. (1990a). Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences* 13: 585-96

———. (1990b). Who is computing with the brain? *Behavioral and Brain Sciences* 13: 632-40

Sigman E., and I. Rock (1974). Stroboscopic movement based on perceptual intelligence. *Perception* 3: 9-28

3

A Feedforward Network for Fast Stereo Vision with Movable Fusion Plane

Paul M. Churchland

Introductory epistemology courses often begin by addressing the problem of "our knowledge of the external world." Typically the problem is posed in the form, "How is one justified in believing a certain class of sentences?" (such as those about the arrangment and character of proximate physical objects). Two major assumptions are thus made surreptitiously central right at the outset. First, to pose the question in this way is to assume that our *representions* of the world are basically sentential or propositional in character. And it is also to assume that the relational features from which an account of *virtuous cognitive activity* is to be drawn must be the various relations that typically connect sentences to each other and to the world – relations such as entailment, coherence, probability, truth and so forth. Thus are we launched on a long tradition of epistemological discussion.

The most recent and in many ways the most useful instance of this tradition is the attempt by classical AI to construct artificial cognitive systems. It is especially useful because, unlike the representational / computational stories earlier composed by philosophers of science and inductive logicians, the stories produced (or reproduced) in AI were quickly implemented on large and powerful machines where their virtues – and their shortcomings – could be made dramatically evident. Where philosophers had spent months or years, pencils in hand, trying to discover failures and counter-examples in some intricate computational scheme, AI researchers could see its wheels fall off much more swiftly, indeed, often in milliseconds. We have learned from the machine-implemented experiments of classical AI, much more firmly than from the scratch-pad experiments of philosophy, how difficult it is to account for the acquisition, administration and deployment of knowledge if we restrict ourselves to the classical representational and computational assumptions of the preceding paragraph.

For this we should be grateful, since a sister discipline invites research down a different path in any case. Empirical neuroscience