OBSERVATION RECONSIDERED*

JERRY FODOR

Department of Psychology
Massachusetts Institute of Technology

Several arguments are considered which purport to demonstrate the impossibility of theory-neutral observation. The most important of these infers the continuity of observation with theory from the presumed continuity of perception with cognition, a doctrine widely espoused in recent cognitive psychology. An alternative psychological account of the relation between cognition and perception is proposed and its epistemological consequences for the observation/theory distinction are then explored.

Granny and I think that things have gone too far, what with relativism, idealism and pragmatism at Harvard, graffiti in the subway stations, and Lord knows what all next. Granny and I have decided to put our foot down and dig our heel in. Granny is particularly aroused about people playing fast and loose with the observation/inference distinction; and when Granny is aroused, she is terrible. "We may not have prayers in the public schools," Granny says, "but by G-d, we will have a distinction between observation and inference."

The observation/inference distinction according to Granny:

"There are", Granny says, "two quite different routes to the fixation of belief. There is, on the one hand, belief fixation directly consequent upon the activation of the senses (belief fixation 'by observation', as I shall say for short) and there is belief fixation via inference from beliefs previously held ('theoretical' inference, as I shall say for short). This taxonomy of the means of belief *fixation* implies, moreover, a corresponding taxonomy of *beliefs*. For, the character of an organism's sensory apparatus—and, more generally, the character of its perceptual psychology—may determine that certain beliefs, if acquired at all, must be inferential and cannot be attained by observation. It is, for example, an accident (of our geography) that our beliefs about Martian fauna are non-observationally acquired. By contrast, it is *not* an accident that our beliefs about the doings of electromagnetic energy in the extreme ultraviolet are all inferential. If there are Martian fauna then were we close enough, we could observe some (unless Martians are *very* small). But making observations in the extreme ultraviolet would require alteration of our sensory/

*Received August 1983; revised October 1983.

Philosophy of Science, 51 (1984) pp. 23–43.

perceptual mechanisms; beliefs about the extreme ultraviolet *must,* for us, all be inferential.

"Some beliefs are thus nonobservational in the nature of things. (To a first approximation, no beliefs are noninferential in the nature of things; any belief *could* be fixed by inference excepting, maybe, tricky ones of the 'I exist' variety.) Moreover, beliefs that are fixed by observation play an interesting and central role in the acquisition of knowledge. (Not, perhaps, so interesting and central as philosophers have sometimes supposed, but still. . . .) For one thing, observationally fixed beliefs tend, by and large, to be more reliable than inferentially fixed beliefs. This is primarily because the etiological route from the fact that *P* to the belief that *P* is metaphorically—and maybe literally—*shorter* in observation than in inference; less is likely to go wrong because there's less that *can* go wrong. And, because beliefs that are fixed by observation tend to be relatively reliable, our rational confidence in our knowledge claims depends very largely on their ability to survive observational assessment.

"Second, the observational fixation of belief plays a special role in the adjudication and resolution of clashes of opinion. When observation is *not* appealed to, attempts to settle disputes often take the form of a search for premises that the disputants share. There is, in general, no point to my convincing you that belief *B* is derivable from theory *T* unless *T* is a theory you endorse; otherwise, my argument will seem to you merely a reductio of its premises. This is a peculiarly nasty property of inferential belief fixation because it means that *the more we disagree about, the harder it will likely be to settle any of our disagreements*. None of this applies, however, when the beliefs at issue are observational. Since observation is not a process in which new beliefs are inferred from old ones, the use of observation to resolve disputes does not depend on a prior consensus as to what premises may be assumed. The moral, children, is approximately Baconian. Don't think; look. Try not to argue."

Also sprach Granny. Recent opinion, however, has tended to ignore these homely truths. In this paper, I want to claim that widely endorsed arguments against the possibility of drawing a principled observation/theory distinction have, in fact, been over-sold. This does not amount quite to Granny's vindication, since I will not attempt to say in any detail what role the notion of observational belief fixation might come to play in a reasonable naturalized epistemology. Suffice it, for present purposes, to have cleared the way for such a reconstruction.

The claim, then, is that there is a class of beliefs that are typically fixed by sensory/perceptual processes, and that the fixation of beliefs in this class is, in a sense that wants spelling out, importantly theory neutral. As a first shot at what the theory neutrality of observation comes to: given the same stimulations, two organisms with the same sensory/perceptual

psychology will quite generally observe the same things, and hence arrive at the same observational beliefs, *however much their theoretical commitments may differ*. This will get some pretty comprehensive refinement as we go along, but it's good enough to start from.

There are, as far as I know, three sorts of arguments that have been alleged to show that no serious observation/inference distinction can be drawn.[1] These are: ordinary language arguments, meaning holism arguments, and de facto psychological arguments. I propose to concentrate, in what follows, mostly on arguments of the third kind; I think that recent changes in the way (some) psychologists view sensory/perceptual processes have significant implications for the present philosophical issues. But it's worth a fast run-through to see why the first two sorts of arguments are also, to put it mildly, less than decisive.

1. The Ordinary Language Argument. The main contention of this paper is that there is a theory-neutral observation/inference distinction; that the boundary between what can be observed and what must be inferred is largely determined by fixed, architectural features of an organism's sensory/perceptual psychology. I'm prepared to concede, however, that this is *not* the doctrine that emerges from attention to the linguistic practices of working scientists. Scientists do have a use for a distinction between what is observed and what is inferred, but the distinction that they have in mind is typically relativized to the inquiry they have in hand. Roughly, so far as I can tell, what a working scientist counts as an experimental *observation* depends on what issue his experiment is designed to settle and what empirical assumptions the design of his experiment takes for granted. One speaks of telescopic observations—and of the telescope as an instrument of observation—because the functioning of the telescope is assumed in experimental designs that give us observations of celestial events. One speaks of observed reaction times because the operation of the clock is assumed in the design of experiments when reaction time is the dependent variable. If, by contrast, it begins to seem that perhaps the clock is broken, it then becomes an issue whether reaction times *are* observed when the experimenter reads the numerals that the clock displays.

That way of using the observation/inference distinction is, of course, responsive to an epistemically important fact: not all the empirical assumptions of an experiment can get tested in the same design; we can't

[1]Well, four really. But I shan't discuss *ontological* approaches which support a distinction between observation terms and others by claiming that only the former denote (eg. because whatever is unobservable is ipso facto fictitious). That the assumptions of the present discussion are fully Realistic with respect to unobservables will become entirely apparent as we proceed.

test all of our beliefs at once. It is perfectly reasonable of working scientists to want to mark the distinction between what's foreground in an experiment and what is merely taken for granted, and it is again perfectly reasonable of them to do so by relativizing the notion of an observation to whatever experimental assumptions are operative. But, of course, if *that* is what one means by the observation/inference distinction, then there is no interesting issue about whether scientific observation can be theory neutral. Patently, on that construal, the theory of the experimental instruments and the (e.g. statistical) theory of the experimental design will be presupposed by the scientist's observational vocabulary, and what the scientist can (be said to) observe will change as these background theories mature. We can *now* observe craters on Venus (small differences in reaction times) because we now have powerful enough telescopes (accurate enough clocks). On this way of drawing it, the observation/inference distinction is inherently heuristic; it is relativized not just to the sensory/perceptual psychology of the observer, but also to the currently available armementarium of scientific theories and gadgets.

Much that is philosophically illuminating can, no doubt, be learned by careful attention to what working scientists use terms like 'observed' and 'inferred' to do; but naturalized epistemology is not, for all that, a merely sociolinguistic discipline. Though one of the things that these terms are used for is to mark a distinction that is beyond doubt theory-relative, that does *not* settle the case against Granny. For, it is open to Granny to argue like this:

"True, there is an epistemologically important distinction that it's reasonable to call 'the' observation/inference distinction, and that is theory-relative. And, also true, it is this theory-relative distinction that scientists usually use the terms 'observed' and 'inferred' to mark. But that is quite compatible with there being another distinction, which it is also reasonable to call 'the' observation/inference distinction, which is also of central significance to the epistemology of science, and which is *not* theory-relative. No linguistic considerations can decide this, and I therefore propose to ignore mere matters of vulgar dialectology henceforth."

In her advanced years, Granny has become quite bitter about ordinary language arguments.

2. *Arguments from Meaning Holism.* Think of a theory (or, mutatis mutandis, the system of beliefs a given person holds) as represented by an infinite, connected graph. The nodes of the graph correspond to the entailments of the theory, and the paths between the nodes correspond to a variety of semantically significant relations that hold among its theorems; inferential relations, evidence relations, and so forth. When the theory is tested, confirmation percolates from node to node along the

connecting paths. When the theory is disturbed—eg. by abandoning a postulate or a principle of inference—the local geometry of the graph is distorted, and the resulting strains are distributed throughout the network, sometimes showing up in unanticipated deformations of the structure of the graph far from the initial locus of the disturbance.

That sort of picture has done a lot of work for philosophers since Quine wrote "Two Dogmas". Most famously, skeptical work. Since—so the story goes—everything connects, the unit of meaning—the minimal context, so to speak, within which the meaning of a theoretical postulate is fixed—appears to be *the whole theory*. It is thus unclear how two theories could dispute the claim that *P* (since the claim that *P* means something different in a theory that entails that *P* than it does in, say, a theory that entails its denial). And, similarly, it is unclear how two belief systems that differ anywhere can help but differ everywhere (since a node is identified by its position in a graph, and since a graph is identified by the totality of its nodes and paths, it appears that only identical graphs can have any nodes in common.)

It is, of course, possible to accept this sort of holism (as, by the way, Granny and I do not) and still acknowledge *some* sort of distinction between observation and inference; eg. the distinction might be construed as epistemic rather than semantic. Suppose every sentence gets its *meaning* from its theoretical context; still, some sentences are closer to the 'edges' of the graph than others, and these might be supposed to depend more directly upon experience for their *confirmation* than sentences further inland do. Quine himself has some such tale to tell. However,—and this is what bears on the present issues—the holism story does suggest that observation couldn't be *theory neutral* in the way that Granny and I think it is. On the holistic account, what you can observe is going to depend comprehensively upon what theories you hold because *what your observation sentences mean depends comprehensively on what theories you hold.*

This is precisely the moral that a number of philosophers have drawn from Quinean holism. For example, here are some quotations from Paul Churchland's recent *Scientific Realism and The Plasticity of Mind*;

> . . . the meaning of the relevant observation terms has nothing to do with the intrinsic qualitative identity of whatever sensations just happen to prompt their non-inferential application in singular empirical judgements. Rather, their position in semantic space appears to be determined by the network of sentences containing them accepted by the speakers who use them (p. 12).
>
> . . . the view that the meaning of our common observation terms is given in, or determined by, sensation must be rejected outright, and

> as we saw, we are left with networks of belief as the bearers or determinants of understanding . . . (p. 13).
>
> . . . a child's initial (stimulus-response) use of, say, 'white' as a response to the familiar kind of sensation, provides that term with no semantic identity. It acquires a semantic identity as, and only as, it comes to figure in a network of beliefs and a correlative pattern of inferences. Depending on what that acquired network happens to be, that term could come to mean *white* or *hot* . . ., or an infinity of other things (14).

And so forth. So Churchland holds, on holistic grounds, that an observation sentence might mean *anything* depending upon theoretical context.

I emphasize that this conclusion is equivalent to the claim that *anything might be an observation sentence* depending upon theoretical context; or, material mode, that *anything might be observed* depending upon theoretical context. For Churchland—as, of course, for many other philosophers—*you can change your observational capacities by changing your theories*. Indeed, Churchland sees in this a program for educational reform. "If our perceptual judgements must be laden with theory in any case, then why not exchange the Neolithic legacy now in use for the conception of reality embodied in modern-era science?" (p. 35). Really well brought up children would not

> . . . sit on the beach and listen to the steady roar of the pounding surf. They sit on the beach and listen to the aperiodic atmospheric compression waves produced as the coherent energy of the ocean waves is audibly redistributed in the chaotic turbulence of the shallows. . . . They do not observe the western sky redden as the Sun sets. They observe the wavelength distribution of incoming solar radiation shift towards the longer wavelengths . . . as the shorter are increasingly scattered away from the lengthening atmospheric path that they must take as terrestrial rotation turns us slowly away from their source. . . . They do not feel common objects grow cooler with the onset of darkness, nor observe the dew forming on every surface. They feel the molecular KE of common aggregates dwindle with the now uncompensated radiation of their energy starwards, and they observe the accretion of reassociated atmospheric H_2O molecules as their KE is lost to the now more quiescent aggregates with which they collide . . . (p. 30).

Oh brave new world/that has such children in it.

Once again: the moral that Churchland (and others) draw from holistic semantic doctrines about beliefs/theories is that an observation sentence can mean anything depending on theoretical context; hence that anything can *be* an observation sentence depending on theoretical context; hence that *there could not be a class of beliefs that must be inferential regardless of what theories the believer espouses*. Churchland's way of putting this is, perhaps, misleading. After all, if the gathering of the dew *is* the accretion of atmospheric H_2O molecules, then of course we do, right now and without technological retraining, observe the accretion of atmospheric H_2O molecules whenever we observe the gathering of the dew; 'observe' is transparent to the substitutivity of identicals. But I don't really think that Churchland (or anybody else party to the present controversy) is seriously confused about this, and I don't propose to carp about it. Indeed, it's easy to fix up. What Churchland must be claiming, on grounds of holism, is that what you can see things *as*—what you can *observe that things are* is comprehensively determined by theoretical context; so that, depending on context, you can, or can learn to, see anything as anything.

Granny and I doubt that you can learn to see anything as anything (that anything can be an observation sentence); but our reasons for doubting this will keep till Section 3. For present purposes, suffice it to repeat the lesson that causal semantic theories have recently been teaching us, viz. that holism may not be true. Specifically, it may not be true that (all) the semantical properties of sentences (/beliefs) are determined by their location in the theoretical networks in which they are embedded; it may be that some of their semantic properties are determined by the character of their attachment to the world (eg., by the character of the causal route from distal objects and events to the tokening of the sentence or the fixation of the belief.) The point is, of course, that their attachment to the world, unlike their inferential role, is something that symbols (/beliefs) can have *severally*; so that, when such attachments are at issue, the morals of holism need not apply.

At a minimum, this suggests a way out of Churchland's dilemma. It will have been clear from the fragments quoted above that Churchland's discussion relies heavily, if implicitly, on the following modus tollens: if the semantics of observation sentences is theory neutral, that must be because observation sentences get their meanings—somehow—from their connections with sensations. But we have good reason to deny that they get their semantics that way. The alternative is that observation sentences get their meanings from their theoretical contexts (from "networks of beliefs").

In fact, however, *neither* of these accounts of the semantics of observation sentences seems particularly attractive, least of all for color terms,

although, as it happens, color terms are Churchland's favorite working examples. It tells against the first alternative that 'white' is typically used to refer to the color of objects, not to the color of sensations; and it tells against the second that the inferential roles of color terms tend to be isomorphic—hence inverted spectrum puzzles—so that color words provide the worst possible cases for 'functional role' theories of meaning. In fact, it looks as though the sensible thing to say about 'white' might be that it means what it does because of the special character of its association (not with a sensation or an inferential role but) *with white things*. To accept that, however, is to reject holism as, anyhow, the *whole* story about the semantics of color terms.

I don't suppose that there's anything much novel in this, and I certainly don't suppose it establishes that there *is* a viable, theory neutral, observation/inference distinction. The point I have been making is merely negative: meaning holism is unequivocally destructive of a theory-neutral notion of observation only if you suppose that *all* the semantic properties of sentences/beliefs are determined by their theoretical context; for, if some are not, then perhaps the essential semantic conditions for being observational can be framed in terms of these. The obvious suggestion would be, on the one hand, that what makes a term observational is that it denotes what is, by independent criteria, an observable property; and, on the other, that what a term denotes is nonholistically (perhaps causally) determined. In light of this, I propose simply not to grant that all the semantic properties of sentences/beliefs are determined by their theoretical context. And Granny proposes not to grant that too.

3. Psychological Arguments. Precisely parallel to the philosophical doctrine that there can be no principled distinction between *observation* and *inference* is the psychological doctrine that there can be no principled distinction between *perception* and *cognition*. The leading idea here is that "perception involves a kind of problem-solving—a kind of intelligence" (Gregory 1970, p. 30). Perception, according to this account, is the process wherein an organism assigns probable distal causes to the proximal stimulations it encounters. What makes the solution of perceptual problems other than mere routine is the fact that, as a matter of principle, any given pattern of proximal stimulation is compatible with a great variety of distal causes; there are, if you like, many possible worlds that would project a given pattern of excitation onto the sensory mechanisms of an organism. To view the mental processes which mediate perception as inferences is thus necessarily to view them as *nondemonstrative* inferences. "We are forced . . . to suppose that perception involves betting on the most probable interpretation of sensory data, in terms of the world

of objects" (Gregory 1970, p. 29). It is worth stressing the putative moral: what mediates perception is an inference from effects to causes. The sort of mentation required for perception is thus not different in *kind*—though no doubt it differs a lot in conscious accessability—from what goes on in Sherlock Holmes' head when he infers the identity of the criminal from a stray cigar band and a hair or two. If what Holmes does deserves to be called cognition, perception deserves to be called cognition too, or so, at least, some psychologists like to say.

Neither Granny nor I have heard of a serious alternative to this view of perception, so let's suppose, for purposes of argument at least, that these psychologists are right. It may then seem that the psychology of perception provides an argument—indeed, quite a direct argument—that observation can't be theory neutral. To see how such an argument might go, consider the following question: if, in general, there are many distal solutions compatible with the perceptual problem that a given sensory pattern poses, how is it possible that perception should ever manage to be univocal (to say nothing of *veridical*)? Why, that is, doesn't the world look to be many ways ambiguous, with one 'reading' of the ambiguity corresponding to each distal layout that is compatible with the current sensory excitation; (as, indeed, a Necker cube *does* look to be several ways ambiguous, with one term of the ambiguity corresponding to each of the possible optical projections from a three dimensional cube onto a two dimensional surface). Assuming, in short, that perception is problem solving, how on earth do perceptual problems ever get solved? As Gregory comments, "it is surely remarkable that out of the infinity of possibilities the perceptual brain generally hits on just about the best one" (1970, p. 29).

All psychological theories that endorse the continuity of perception with problem solving offer much the same answer to this question: viz. that though perceptual analyses are underdetermined by sensory arrays, it does not follow that they are underdetermined *tout court*. For, perceptual analyses are constrained not just by the available sensory information, but also by such prior knowledge as the perceiver may bring to the task. What happens in perceptual processing, according to this account, is that sensory information is interpreted by reference to the perceiver's background theories, the latter serving, in effect, to rule out certain etiologies as implausible causal histories for the present sensory array. Only thus is sensory ambiguity resolved; and, if perception is typically veridical, that's because the background theories that organisms exploit in perceptual analysis are, for the most part, true.

Accepting this account of the perceptual reduction of sensory ambiguity is, of course, fully compatible with stressing the analogy between per-

ception and problem solving. There are many, many ways that the hairs and the cigar band could have come to where Holmes found them; many projections, if you like, of possible criminals onto actual clues. How, then, is it possible—even in principle—that Holmes should solve the crime? Answer: Holmes knows about the clues, but he knows a lot more too; and his background knowledge comes into play when the clues get unravelled. Jones couldn't have left brown hairs because Jones is blond; Smith couldn't have left the cigar band because he only smokes iced tea. Bentley, however, has brown hair and his dog collects cigar bands; so Bentley and his dog it must have been. The clues underdetermine the criminal, but the clues plus background knowledge may be univocal up to a very high order of probability. The trick—the trick that problem solving *always* amounts to—is having the right background information and knowing when and how to apply it. So too in the case of perception, according to the cognitivists.

What has all this to do with reconsidering observation? The point is that, if the present story is right, then the appeal to a background theory is *inherent* in the process of perceptual analysis. Perception wouldn't work without it because the perceptual problem is the reduction of sensory ambiguity, and that problem is solved only when one's sensory information is interpreted in the light of one's prior beliefs. So, the one thing that perception *couldn't* be, on this account of how it works, is theory neutral. Indeed, this is precisely the moral that a number of philosophers have drawn from the psychological texts. Thus, Thomas Kuhn remarks that "the rich experimental literature [in psychology] . . . makes one suspect that something like a paradigm is prerequisite to perception itself. What a man sees depends both upon what he looks at and also upon what his previous visual-conceptual experience has taught him to see" (Kuhn 1962, p. 113). Kuhn clearly thinks that, among the "visual-conceptual experiences" that can work such alterations in perception is the assimilation of scientific doctrine: ". . . paradigm changes do cause scientists to see the world of their research-engagements differently. . . . It is as elementary prototypes for these transformations of the scientist's world view that the familiar demonstrations of a switch in visual gestalt prove so suggestive" (1962, p. 111). Nelson Goodman reads the experimental literature on perception in much the same way. "That we find what we are prepared to find (what we look for or what forcefully affronts our expectations), and that we are likely to be blind to what neither helps nor hinders our pursuits, are commonplaces . . . amply attested in the psychological laboratory. [See also Goodman's *Languages of Art,* where this view of perceptual psychology is strikingly in evidence.]" (Goodman 1978, p. 14).

In fact, however, it is unclear that that's what the psychological lab-

oratory *does* attest, and thereby hangs a puzzle. For if we ought to be impressed by the degree to which perception is interpretive, contextually sensitive, labile, responsive to background knowledge and all that, we surely ought also to be impressed by the degree to which it is often bull headed and recalcitrant. In fact, many of the standard psychological demonstrations seem to point both morals at the same time. Consider the famous Muller-Lyre figures. The text-book story goes like this:

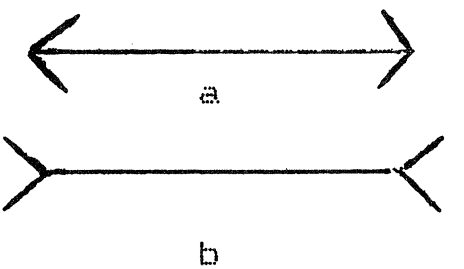

Figure 1. The Muller-Lyre Illusion

when the arrow heads bend in, as in 1a, the figure is unconsciously interpreted in three dimensional projection as a convex corner with its edge emerging towards the viewer from the picture plane. Conversely, when the arrow heads bend out, as in 1b, the figure is unconsciously interpreted in three dimensional projection as a concave corner with its edge receding from the viewer. It follows that the center line is interpreted as *further from the observer* in 1b than in 1a. Since, however, the two center lines are in fact of the same length, their retinal projections are identical in size. This identity of retinal projection could be compatible with the three dimensional interpretation of the figures only if the center line is longer in figure b than in figure a; two objects at different distances can have the same retinal projection only if the more distant object is larger. So size constancy operates (to compensate, as one might say, for what appears to be the apparent effect of distance) and the two lines are perceived as differing in length. See what a nice regard for consistency the unconscious has, Freud to the contrary notwithstanding. There is abundant empirical evidence for this explanation including, notably, the fact that children, having had less experience with edges and corners than adults, are correspondingly less susceptible to the illusion.

The Muller-Lyre illusion thus appears to be, and is often cited as, a prime example of how background information—in this case a complex of assumptions about the relations between three-dimensional objects and

their two-dimensional projections—can affect the perceptual analysis of a sensory array. 'What', one might ask, 'could be clearer evidence of the penetration of perception by information that is *not* available at the retina?' On the other hand, there's this: The Muller-Lyre is a *familiar* illusion; the news has pretty well gotten around by now. So, it's part of the 'background theory' of anybody who lives in this culture and is at all into pop psychology that displays like Figure 1 are in fact misleading and that it always turns out, on measurement, that the center lines of the arrows are the same length. Query: *Why isn't perception penetrated by THAT piece of background theory?* Why, that is, doesn't *knowing* that the lines are the same length make it *look as though* the lines are the same length? (For that matter, since one knows perfectly well that Figure 1 is a drawing in two dimensions, why doesn't *that* information penetrate perception, thereby blocking the three dimensional interpretation and cancelling the illusion?) This sort of consideration doesn't make it seem at all as though perception is, as it's often said to be, saturated with cognition through and through. On the contrary, it suggests just the reverse: that how the world looks can be peculiarly unaffected by how one knows it to be. I pause to emphasize that the Muller-Lyre is by no means atypical in this respect. To the best of my knowledge, all the standard perceptual illusions exhibit this curiously refractory character: knowing that they *are* illusions doesn't make them go away.[2]

I hope that the polemical situation is beginning to seem a little queer. On the one hand, reflection upon the impoverishment and ambiguity of sensory information leads, by a plausible route, to the analysis of perception as a form of problem solving in which proximal stimulations are interpreted in light of some background theory accessible to the perceiver. This makes it seem that how the world is perceived to be ought to depend very largely on the perceiver's prior beliefs and expectations; hence the perceptual effects of cognitive set that psychologists of the 'New Look' persuasion made a living by advertizing. But, on the other hand, there are these curious and persuasive perceptual *im*plasticities, cases where knowing doesn't help seeing. It is, of course, reflection on examples of

[2]Interestingly enough, Jerome Bruner, in his foundational 'New Look' disquisition "On Perceptual Readiness", takes note of this point using, in fact, the same examples I have cited. But he makes nothing of it, remarking only that the persistence of illusions in face of contrary background knowledge, though it militates against the "utter indistinguishability of perceptual and more conceptual inferences . . . must not lead us to overlook the common feature of inference underlying so much of cognitive activity" (1973, p. 8). The issue, however, is not whether some inferences are "more conceptual" than others—whatever, precisely, that might mean—or even whether perception is in some important sense inferential. What's at issue is rather: how much of what you know actually does affect the way you see. Failing to distinguish among these questions was, in my view, the original sin of New Look psychological theorizing.

the second sort that keeps Granny going. These are the cases where the idea of theory-neutral observation can get a toe-hold. The problem is: which sorts of cases ought we to believe? And, while we're at it, *how can a theory of perception accommodate the existence of both?*

We come to the main point at last. The New Look idea that perception is a kind of problem solving does not, all by itself, imply the theory dependence of observation. Philosophers who read that moral in the psychological texts read the texts too fast. (Granny says that a little psychology is a dangerous thing and inclineth a man to relativism.) To get from a cognitivist interpretation of perception to any epistemologically interesting version of the conclusion that observation is theory dependent, you need *not only the premise that perception is problem solving, but also the premise that perceptual problem solving has access to ALL (or, anyhow, arbitrarily much) of the background information at the perceiver's disposal*. Perceptual implasticities of the sorts we've just been noticing make it highly implausible, however, that this second premise is true.

All this suggests that we'd better distinguish between two questions that up 'til now we've been treating as the same: the question whether perception is a kind of problem solving (i.e. whether observation is inferential) and the question whether perception is comprehensively penetrated by background beliefs (i.e. whether observation can be theory-neutral). It is entirely possible—to put the point another way—to steer a middle course between Granny and Jerome Bruner: to agree with Bruner (as against Granny) that there is an important sense in which observation is a kind of inference, but also to agree with Granny (as against Harvard relativists) that there is, in perception, a radical isolation of how things look from the effects of much of what one believes.

Since it is the second issue, rather than the first, that raises all the epistemological questions, this seems to be a moral victory for Granny. If, for example, the inferential character of perception is, as I'm supposing, compatible with the theory neutrality of observation, then *nothing* follows from perceptual psychology about whether scientists who accept radically different theories can observe the same phenomena. In particular, on this view, it would *not* follow from the inferential character of perception that "the infant and the layman . . . cannot see what the physicist sees" (Hanson 1961, p. 17), or that "[when the physicist looks at an X-ray tube] . . . he sees the instrument in terms of electrical circuit theory, thermodynamic theory, the theories of metal and glass structure, thermionic emission, optical transmission, refraction, diffraction, atomic theory, quantam theory and special relativity" (pp. 15–16). Similarly, on this account, the inferential character of perception leaves it open that the children whom Churchland wants to teach not to see the gathering of the

dew might, thank God, see things much the same way after they've learned physics as they did before. The argument for the relativity of observation requires, to repeat, not just the inferential character of perception, but the idea that *all* your background knowledge, including especially your scientific theories, is accessible as premises for perceptual integration. By contrast, if you think that perception, though inferential, is nevertheless encapsulated from much of what the perceiver believes, the common epistemic situation of the scientist and the layman starts to show through. There is, perhaps, just one perceptual world, though the experts sometimes know more about it than the amateurs.

What might the psychology of perception look like if observation is *both* inferential and theory neutral? I'll say a word about this before returning to the epistemological issues.

The view that perception is problem solving, though it takes the distinction between perception and *cognition* as heuristic, takes quite seriously the distinction between perception and *sensation*. Sensory processes, according to this account, merely register such proximal stimulations as an organism's environment affords. It's left to cognitive processes—notably the perceptual ones—to interpret sensory states by assigning probable distal causes. So we have the following picture: sensation is responsive solely to the character of proximal stimulation and is noninferential. Perception is both inferential and responsive to the perceiver's background theories. It is not, of course, an accident that things are supposed to line up this way; inference requires premises. Perceptual processes *can* be inferential because the perceiver's background theory supplies the premises that the inferences run on. Sensory processes *can't* be inferential because they have, by assumption, no access to the background theories in light of which the distal causes of proximal stimulations are inferred. The moral is that, if you want to split the difference between Granny and the New Look, you need to postulate a tertium quid; a kind of psychological mechanism which is both encapsulated (like sensation) and inferential (like cognition). The apparent contradiction between inference and encapsulation is resolved by assuming that the access to background theory that such mechanisms have is sharply delimited; indeed, delimited by the intrinsic character of the mechanisms.

I won't say much about this here since I've set out the psychological story at some length in a previous study (see Fodor 1983) and I'm anxious to return to the philosophical morals. Suffice it just to suggest, by way of a brief example, what the organization of such "modular" perceptual mechanisms might be like.

It's plausible to assume that the perceptual analysis of speech typically effects an assignment of sentence tokens to sentence types. One reason it's plausible to assume this is that it's obviously true. Another reason is

that understanding what someone says typically requires knowing what form of words he uttered, and to assign an utterance to a form of words *is* to assign a token to a type. Cognitive psychology proceeds by diagnosing functions and postulating mechanisms to perform them; so let's assume that there is some psychological mechanism—a *parser,* let's call it—whose function is this: it takes sensory (as it might be, acoustic) representations of utterances as inputs and produces representations of sentence types (as it might be, linguistic structural descriptions) as outputs. No doubt this way of setting up the problem assumes a lot that a lot of you won't want to grant—for example, that there are psychological mechanisms, and that they are properly viewed as functions from one sort of representations onto another. However, remember the context: we've been wondering what current psychological theory implies about the observation/inference distinction. And the sort of psychological theory that's current is the one I've just outlined.

There is abundant empirical evidence—with which, however, I won't bother you—that parsing has all the properties that make psychologists want to say that perception is inferential. All the indications are that the acoustic character of an utterance significantly underdetermines its structural description, so the parser—if it is to succeed in its function—will have to know a lot of background theory. This isn't, by the way, particularly mysterious. Consider the property of being a noun—a sort of property that some utterances surely have and that adequate structural descriptions of utterances must surely mark. Patently, that property has no sensory/acoustic correspondent; there's nothing that nouns qua nouns sound like, or look like on an oscilliscope. So a mechanism that can recognize utterances of nouns as such must know about something more than the acoustic/sensory properties of the tokens it classifies; in this case, something about the language that it parses; i.e. it has to know which words in the language are nouns.

Well, then, what would it be like for the parser to be a module? A simple story might go like this; a parser for L contains a grammar of L. What it does when it does its thing is: it infers from certain acoustic properties of a token to a characterization of certain of the distal causes of the token (eg. to the speaker's intention that the utterance should be a token of a certain linguistic type). Premises of this inference can include: whatever information about the acoustics of the token the mechanisms of sensory transduction provide, whatever information about the linguistic types in L the internally represented grammar provides, *and nothing else*. It is, of course, the closure condition that makes the parser modular.

Compare a New Look parser. In the extreme case, a New Look parser can bring to the process of assigning structural descriptions *anything that*

the organism knows (or believes, or hopes, or expects . . . etc). For example, a New Look parser knows how *very* unlikely it is that anyone would say, right smack in the course of a philosophical lecture on observation and inference: "Piglet gave Pooh a stiffening sort of nudge, and Pooh, who felt more and more that he was somewhere else, got up slowly and began to look for himself." So if someone *were* to say that, right smack in the middle of a philosophical lecture on observation and inference, a New Look parser would presumably have a lot of trouble understanding it; by definition, a New Look parser tends to hear just what it expects to hear. By the way, this example suggests one of the reasons why encapsulated perceptual modules might be quite a good thing for an organism to have: Background beliefs, and the expectations that they engender, from time to time prove *not to be true*. That doesn't matter so much when they are background beliefs about observation and inference, or about Pooh and Piglet. When, however, they are background beliefs about Tigger, it's a different story. Tiggers bounce. And bite.

I won't try to convince you that the parser—or any other perceptual mechanism—actually *is* modular; what I want to urge, for present purposes, is just that *if* perception is modular (inferential but encapsulated), then that has serious implications for the putative psychological arguments against the theory neutrality of observation. I have a scattering of points to make about this.

First, and most important, if perceptual processes are modular, then, by definition, bodies of theory that are inaccessible to the modules *do not affect the way the perceiver sees the world*. Specifically, perceivers who differ profoundly in their background theories—scientists with quite different axes to grind, for example—might nevertheless see the world in *exactly* the same way, so long as the bodies of theory that they disagree about are inaccessible to their perceptual mechanisms.

Second, the modularity story suggests not only that something can be made of the notion of theory neutral observation, but also that something can be made of the notion of observation *language*; i.e. that—much current opinion to the contrary notwithstanding—there is a good sense in which some terms (like 'red', as it might be) are observational and others (like 'proton', as it might be) are not. Suppose that perceptual mechanisms are modular and that the body of background theory accessible to processes of perceptual integration is therefore rigidly fixed. By hypothesis, only those properties of the distal stimulus count as observable which terms in the *accessible* background theory denote. The point is, no doubt, entirely empirical, but I am willing to bet lots that 'red' will prove to be observational by this criterion and that 'proton' will not. This is, of course, just a way of betting that Hanson, Kuhn, Churchland, Goodman and Co. are wrong; that physics doesn't belong to the accessible background.

There are other, more exciting cases where we are already in a pretty good position to say which properties of distal objects will count as observable, hence which terms will count as observation vocabulary. The case of parsing is among these. This is because it is plausible to suppose that the background theory accessible to a modularized parser would have to be a grammar, and we know, more or less, what sorts of properties of sentences grammatical descriptions specify. So then, applying the present criterion to the present assumptions, the observable linguistic properties of utterances of sentences ought to include things like: being an utterance of a sentence, being an utterance of a sentence that contains the word 'the', being an utterance of a sentence that contains a word that refers to trees . . . and so forth, depending on details of your views about what properties of sentences linguistic structural descriptions specify. By contrast, what would *not* count as observable on the current assumptions are such properties of sentences as: being uttered with the intention of deceiving John; being ill-advised in the context, containing a word that is frequently used in restaurants where they sell hamburgers . . . and so forth. It should be noted in passing that this sort of account permits one to distinguish sharply between observable properties and *sensory* properties. If sensory properties are ones that *non*inferential psychological mechanisms respond to, then the sensory properties of utterances are plausibly all acoustic and almost all inaccessible to consciousness.

Third point: what I've been saying about modularity so far is equivalent to the claim that perceptual processes are 'synchronically' impenetrable by—insensitive to—much of the perceiver's background knowledge. Your current sophistication about the Muller-Lyre is inaccessible to the module that mediates visual form perception and does not, therefore, serve to dispel the illusion. But this leaves open the question whether perception may be 'diachronically' penetrable; in effect, whether experience and training can affect the accessability of background theory to perceptual mechanisms.

To deny diachronic penetrability would be to claim, in effect, that *all* the background information that is accessible to modular perceptual systems is endogenously specified, and that is viewed as implausible even by mad dog nativists like me. For example, parsing may be modular, but children must learn *something* about their language from the language that they hear; why else would children living in China so often grow up speaking Chinese? The point about the diachronic penetrability of perception is, however, just like the point about its synchronic penetrability: it offers an argument for the continuity of perception with cognition only if just any old learning or experience can affect the way you see, and there is no reason at all to suppose that that is so. Perhaps, on the contrary, perception is diachronically penetrable only within strictly—maybe

endogenously—defined limits. Not only do your current Copernican prejudices fail to much dispel the apparent motion of the Sun, it may be that there is *no* educational program that would do the trick; because it may be that the inaccessibility of astronomical background to the processes of visual perceptual integration is a consequence of innate and unalterable architectural features of our mental structure. In this case, our agreement on the general character of the perceptual world might transcend the particularities of our training and go as deep as our common humanity. Granny and I hope that this is so since common humanity is something that we favor.

I return now to more strictly epistemological concerns. Two points and I'll have done.

First, if Granny wants to appeal to modularity psychology as a way of holding onto theory-neutral observation, she is going to have to give a bit. In particular, she is going to have to distinguish between *observation* and *the perceptual fixation of belief*. It is only for the former that claims for theory neutrality have any plausibility.

Thus far, I've been emphasizing that psychological sophistication doesn't change the way the Muller-Lyre *looks*. Knowing that it's an illusion—even knowing how the illusion works—doesn't make the effect go away. But if one side of perception is about the look of things, the other side is about how things are judged to be; and it bears emphasis that how the Muller-Lyre looks doesn't, in the case of a sophisticated audience, much affect the perceptual beliefs that its observers come to have. I assume, for example, that you're not remotely tempted to suppose that the center line in figure b actually is longer than the center line in figure a; and the reason you're not is that *the mechanisms of belief fixation, in contrast to the presumptive perceptual modules, ARE in contact with background theory*. Belief fixation, unlike the fixation of appearances—what I'm calling observation—is a *conservative* process; to a first approximation, it uses everything you know.

Here is one way to conceptualize the situation: the fixation of perceptual belief effects a reconciliation between the character of current sensory stimulation, as analyzed by modular processors, and background theory. The modular systems might be thought of as proposing hypotheses about the distal sources of sensory stimulation; these hypotheses are couched in a restricted (viz. observational) vocabulary and are predicated on a correspondingly restricted body of information: viz. current sensory information together with whatever fragment of background theory the modules have access to. The hypotheses that modular systems propose are then compared with the rest of the organism's background theory, and the perceptual fixation of belief is consequent upon this comparison.

So, to a first approximation, the activity of the modules determines

what you would believe if you were going on the appearances alone. But, of course, this is *only* a first approximation since, as remarked above, modules deal not only in a restricted body of background knowledge, but also in a restricted conceptual repertoire. There are some hypotheses that modules *never* offer because they have no access to a vocabulary in which to express them: hypotheses about the instantiation of nonobservable properties such as, for example, that what's currently on view is a proton. So one might better put it that the activity of modules determines what you would believe *about the appearances* if you were going just on the appearances. Less gnomically: modules offer hypotheses about the instantiation of observable properties of things, and the fixation of *perceptual belief* is the evaluation of such hypotheses in light of the totality of background theory. According to this usage, what you observe is related to what you believe in, something like the way that what you want is related to what you want on balance.

It should be clear from all this that even if Granny gets the theory-neutrality of observation, she is unlikely to get anything remotely like its infallibility. For starters, only a faculty of belief fixation can be infallible and, according to the present story, the psychological mechanisms that are informationally encapsulated do not, in and of themselves, effect the fixation of belief. Anyhow—beside this somewhat legalistic consideration—the infallibility of observation would presumably require the introspective availability of its deliverances; and, though I suppose one usually knows how things look to one, it seems to be empirically false that one always does. If, for example, the story I told about the Muller-Lyre is true, then the existence of the illusion turns on the fact that one sees the figures as three dimensional corners. But it is *not* introspectively obvious that one sees them that way, and the psychologists who figured out the illusion did so not by introspecting but by the usual route of theory construction and experimentation. (Similarly, a crucial issue in the history of the psychology of color perception was whether yellow looks to be a mixed hue. It is *now*—post-theoretically—introspectively obvious that it does not.)

'But look', you might say, growing by now understandably impatient, 'if the notion of observation we're left with is as attenuated as it now appears to be, what, epistemologically speaking, is it good for? Haven't you and your Granny really given away everything that the opposition ever wanted?'

I quote from Norwood Russell Hanson: "To say that Tycho and Kepler, Simplicius and Galileo, Hooke and Newton, Priestly and Lavoisier, Soddy and Einstein, De Broglie and Born, Heisenberg and Bohm all make the same observations but use them differently is too easy. This parallels the too-easy epistemological doctrine that all normal observers see the same

things in x, but interpret them differently. It does not explain controversy in research science" (Hanson 1961, p. 13). (In Hanson's text, the second sentence appears as a footnote at the point where I have inserted it.) Now, on the view of science that Granny and I hold to, this is worse than the wrong answer; it's the answer to the wrong question. It is no particular puzzle, given the nondemonstrative character of empirical inference, that there should be scientific controversy. Rather, as the skeptical tradition in philosophy has made crystal clear, the epistemological problem *par excellence* is *to explain scientific consensus*; to explain how it is possible, given the vast and notorious underdetermination of theory by data, that scientists should agree about so much so much of the time.

What Granny and I think is that part of the story about scientific consensus turns crucially on the theory-neutrality of observation. Because the way one sees the world is largely independent of one's theoretical attachments, it is possible to see that the predictions—even of theories that one likes a lot—aren't coming out. Because the way one sees the world is largely independent of one's theoretical attachments, it is often possible for scientists whose theoretical attachments differ to agree on what experiments would be relevant to deciding between their views, and to agree on how to describe the outcomes of the experiments once they've been run. We admit, Granny and I do, that working scientists indulge in every conceivable form of fudging, smoothing over, brow beating, false advertising, self-deception, and outright rat painting—all the intellectual ills that flesh is heir to. It is, indeed, a main moral of this paper that, in many important ways, *scientists are a lot like us*. Nevertheless, it is perfectly obviously true that scientific observations often turn up unexpected and unwelcome facts, that experiments often fail *and are often seen to do so,* in short that what scientists observe isn't determined solely, or even largely, by the theories that they endorse, still less by the hopes that they cherish. It's *these* facts that the theory neutrality of observation allows us to explain.

The thing is: if you don't think that theory neutral observation can settle scientific disputes, you're likely to think that they are settled by appeals to coherence, or convention or—worse yet—by mere consensus. And Granny—who is a Realist down to her tennis sneakers—doesn't see how any of those could compel *rational* belief. Granny and I have become pretty hardened, in our respective old ages; but we're both still moved by the idea that belief in the best science is rational because it is objective, and that it is objective because the predictions of our best theories can be *observed to be true*. I'm less adamant than Granny is, but I don't find the arguments against the theory neutrality of observation persuasive, and I think that the theory neutrality of observation is a doctrine that Realists have got to hold onto. "Help stamp out creeping pluralism", Granny says;

"give 'em an inch and they'll take a mile!" "Right on (with certain significant qualifications)!" say I.

REFERENCES

Bruner, J. (1973), "On Perceptual Readiness", in Anglin, J. (ed.), *Beyond The Information Given*. New York: W. W. Norton & Co.

Churchland, P. (1979), *Scientific Realism and The Plasticity of Mind*. Cambridge, England: Cambridge University Press.

Fodor, J. (1983), *The Modularity of Mind*. Cambridge, Mass: M.I.T. Press.

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

Gregory, R. (1970), *The Intelligent Eye*. New York, NY: McGraw-Hill Book Company.

Hanson, N. R. (1961), *Patterns of Discovery*. Cambridge, England: Cambridge University Press.

Kuhn, T. (1962), *The Structure of Scientific Revolutions*. Chicago, Illinois: The University of Chicago Press.