Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor

Paul M. Churchland

Philosophy of Science, Vol. 55, No. 2 (Jun., 1988), 167-187.

Stable URL:
http://links.jstor.org/sici?sici=0031-8248%28198806%2955%3A2%3C167%3APPATNA%3E2.0.CO%3B2-8

Philosophy of Science is currently published by The University of Chicago Press.

Philosophy of Science

June, 1988

PERCEPTUAL PLASTICITY AND THEORETICAL NEUTRALITY: A REPLY TO JERRY FODOR*

PAUL M. CHURCHLAND†

Department of Philosophy
University of California, San Diego

The doctrine that the character of our perceptual knowledge is plastic, and can vary substantially with the theories embraced by the perceiver, has been criticized in a recent paper by Fodor. His arguments are based on certain experimental facts and theoretical approaches in cognitive psychology. My aim in this paper is threefold: (1) to show that Fodor's views on the impenetrability of perceptual processing do not secure a theory-neutral foundation for knowledge; (2) to show that his views on impenetrability are almost certainly false; and (3) to provide some additional arguments for, and illustrations of, the theoretical character of all observation judgments.

The idea that observational knowledge always and inevitably involves some theoretical presuppositions or prejudicial processing is an idea that has provoked much discussion in recent years, for its consequences are profound. If observation cannot provide a theory-neutral access to at least some aspects of reality, then our overall epistemic adventure contains both greater peril, and greater promise, than we might have thought. The first and perhaps the most important consequence is that we must direct our attention away from foundational epistemologies, and toward epistemologies that tell a more global story of the nature of theoretical justification and rational belief. A second consequence is that our current observational ontology is just one such ontology out of an indefinitely large number of alternative observational ontologies equally compatible

*Received February 1987; revised April 1987.

†Thanks to Patricia Churchland, Michael Stack, Stephen Stich, Philip Kitcher, Patricia Kitcher, and two anonymous referees for useful criticism and discussion of earlier drafts.

Philosophy of Science, 55 (1988) pp. 167–187.

with our native sensory apparatus. And a third consequence is that, since some theoretical frameworks are markedly superior to others, the quality of our observational knowledge is, in principle, improvable. If the conceptual framework in which our perceptual responses to the world are habitually framed were to be replaced by a more accurate and penetrating conception of physical reality, then our newly framed perceptual judgments could be significantly more revealing of the structural properties and the dynamical details of our perceptual environment.

The motivation for such a view is not purely philosophical. Perceptual psychology provides supporting evidence in the form of experiments designed to illustrate both the inevitable ambiguity of perceptual situations and the cunning resolution of those ambiguities at the hands of general assumptions imposed by 'higher' cognitive centers (Gregory 1970, 1974; Bruner 1973; Rock 1983). These 'New Look' ideas, however, have recently come under interesting attack from within cognitive and computational psychology itself. The complaint is that these ideas have exaggerated the extent to which perceptual processing is under the control of the higher cognitive centers. And the counterclaim is that the job of reducing ambiguity is conducted largely or entirely by peripheral 'modules' whose activities are insulated from, and quite insensitive to, the fickle content of human belief.

It is here that Jerry Fodor enters the debate. In a recent paper in this journal (Fodor 1984), he marshalls the alleged modularity of our perceptual systems in criticism of various claims made by Hanson (1958), Kuhn (1962), Churchland (1979), and others, concerning the theory-laden character of perceptual knowledge and the holistic nature of the human epistemic enterprise. My principal aim in this paper is to show that Fodor's specific claims about the psychology of human perception are mostly irrelevant to the epistemological issues at stake here. His discussion serves more to muddy the waters than to clarify them, for even if the modularity/encapsulation thesis is correct—which almost certainly it is not—it contains no significant message concerning the traditional epistemological issues. It is, in short, a red herring. In what follows, I shall try to defend and expand on the specific claims listed in my opening paragraph, against the several criticisms directed at them in Fodor's paper.

There are three principal ways in which any perceptual belief may fail of theoretical neutrality: in its causal history or *etiology,* in its *semantics,* and in the purely *extensional structure* of the ontology it presupposes. In his (1984), Fodor has much to say on the first topic, a little on the second, and he does not discuss the third. Since he does not address what I have called "extensional bias" (Churchland 1975), and space does not permit its exploration here, I shall merely emphasize its existence, and move on. What follows will be focused on the first two loci of epistemic prejudice.

1. The Etiology of Perceptual Belief

1.1. Does Encapsulated Processing Buy Us Theory-Neutral Perceptions? I shall pass over Fodor's opening discussion in order to address immediately what he describes as his main point (p. 35). Fodor, of course, is quite aware that early perceptual processing very likely does involve many elements that resemble or correspond to general empirical 'assumptions' about the world (for example, the three-dimensionality of space, the spatial and temporal continuity of common objects, the sharp change of luminance at a body's boundaries, color constancy through changing environments, the occlusion of distant bodies by proximate ones, etc.), and to 'inferences' drawn or 'hypothesis selections' made in accordance with a system of such default assumptions. On this view, the etiology of perceptual beliefs looks highly, even dramatically, theoretical in character, as Fodor himself remarks (p. 34).

But Fodor's view, to a first approximation at least, is that (a) the assumptions involved in early processing are endogenously fixed in all of us, and (b) the processing in which they play a role is insulated from any contrary assumptions or theories—indeed, from any *additional* assumptions whatever—that the perceiver may subsequently come to believe. Our perceptual processing is thus encapsulated; it delivers outputs to the higher cognitive centers, but it is impenetrable to any inputs from them. The result, according to Fodor, is that all humans are fated to share a common perceptual experience, an experience whose character is not subject to change as a function of any theories we may come to embrace. There is therefore an important sense, he concludes, in which human perception is neutral vis-à-vis the rough and tumble of competing theories. There is an unchanging perspective, on at least some parts of reality, that all human theorists must share in common.

The evidence in support of these claims is twofold. First, Fodor cites a number of experimental facts that illustrate, not the plasticity of perception, but rather the occasional rigidity of our perceptual deliverances (for example, the persistence of certain illusions such as the Müller-Lyer illusion) even in cases where we know them to be mistaken. Second, he claims that if perception is to be theory dependent in any epistemologically interesting sense, then the perceptual modules must have *"access to ALL (or anyhow, arbitrarily much) of the background information at the perceiver's disposal"* (p. 35). Given the rigidity just cited, however, he concludes that the modules at issue lack such access, and hence that perception is not theory dependent in any interesting sense.

Let us suppose, for the moment, that our perceptual modules are indeed informationally isolated in the fashion claimed. That is to say, they embody a systematic set of endogenous or genetically implanted assumptions

about the world, whose influence on perceptual processing is unaffected by any additional or contrary information.

Now this may be a recipe for a certain limited *consensus* among human perceivers, but it is hardly a recipe for theoretical *neutrality,* and it is plain misleading to use this latter term to describe what encapsulation might secure. As conceived within the relevant dialectical tradition, an observation judgment is *theory neutral* just in case its truth is not contingent upon the truth of any general empirical assumptions, just in case it is free of potentially problematic presuppositions. If an observation judgment does have such presuppositions, its theory-laden character will in no way be reduced by hard-wiring those presuppositions into the process by which the judgment is produced, and by closing the process to all contrary information.

If everyone is a hopeless slave of the same hard-wired theory, then what we have is a universal dogmatism, not an innocent Eden of objectivity and neutrality. The alleged cognitive impenetrability of our perceptual processing does nothing to reduce the extent to which the truth of our perceptual beliefs is contingent upon the truth of those background empirical assumptions or theories in which they are semantically embedded. Encapsulation does nothing to ensure the truth of our perceptual beliefs, not even their 'truth in general' or their 'truth under normal circumstances'. Nor does it ensure their epistemological integrity relative to competing interpretations of our sensory input. It merely dooms us to a single point of view, a point of view that is epistemologically just as problematic as any of the infinity of other sets of empirical assumptions that might have been hard-wired into us instead.

Fodor's premises, therefore, do not buy him anything like the theoretical neutrality of our perceptual judgments. An unchangable set of prejudicial empirical assumptions is still a set of prejudicial empirical assumptions.

Fodor's premises may seem to solve, at least, the problem of incommensurability, by guaranteeing some effective communication, at the observation level, between ideologically diverse human theorists. But as we shall see at the end of this section, they fail to guarantee this also, since a rigidity in our early perceptual processing is entirely consistent with plasticity at the level of conceptual apprehension and discursive judgment. And despite a popular misconception on this point, communication was never the real problem anyway. The epistemological problem of incommensurable alternatives arises most clearly and forcefully within a *single* individual, one who is 'bitheoretical'. Putting Fodor aside for a moment, consider someone who has internalized two competing theories, and has learned two correspondingly different ways of perceiving the relevant aspects of the world, but is torn over which of these two global

packages to choose. It is not communication that is the problem here (he can perfectly well understand himself); the problem is theoretical evaluation and rational choice in the absence of a neutral touchstone.

I am not at this point arguing that Fodor's encapsulation thesis is false: only that it would not secure for us any theory-neutral foundation for knowledge, even were it true. Fodor's hard-wired consensus is a sham neutrality: it mistakes the presumed universality of our prejudice for the absence of any prejudice. And hard-wired or no, that consensus would last only until the first mutant or alien comes along, to confront us with a different perceptual point of view.

In fact, we begin to become such mutants or aliens ourselves, when we change our sensory modalities by augmenting them with unusual instruments such as phase-contrast microscopes, deep-sky telescopes, long-baseline stereoscopes, infrared scopes, and so forth. And the metamorphosis is completed when, after years of professional or amateur practice, we learn to see the world appropriately and efficiently with these new senses. This learning requires both that we suppress certain habits of processing 'natural' to the naked eye and to the familiar world of middle-sized material objects, and that we learn to process the retinal data in novel ways that are appropriate to the unfamiliar features one perceives by these novel means (for example, interference patterns, diffraction rings, dark nebulae, fusion planes, temperature gradients, etc.). Reflections such as these do begin to challenge Fodor's factual claim of encapsulation or impenetrability. Let us therefore focus on the evidence he cites in support of that claim.

1.2. Is the Impenetrability Thesis Correct? Visual illusions are good illustrations of the assumptions involved in early processing, since the illusion is often the result of the persistent operation of some assumption that is appropriate for most situations, but which is inappropriate for the particular situation at issue. Fodor cites the stubborn persistence of various visual illusions, even when we know that we are being misled, and even when we have the information about the inappropriate assumptions responsible for the illusory experience. Why, Fodor asks, doesn't *this* information affect the way we see the world, and thereby undo the illusion? His answer is that our perceptual processing is guided by mechanisms or assumptions that cannot be successfully overridden by contrary assumptions imposed from the outside.

A first response is to point out the great many illusions and visual effects whose character shows that our visual modules are indeed penetrable by higher cognitive assumptions. Consider the wide range of ambiguous figures, such as the Duck/Rabbit, the Old/Young Woman, the Necker cube, or the Vase/Faces. Such examples are ambiguous w/

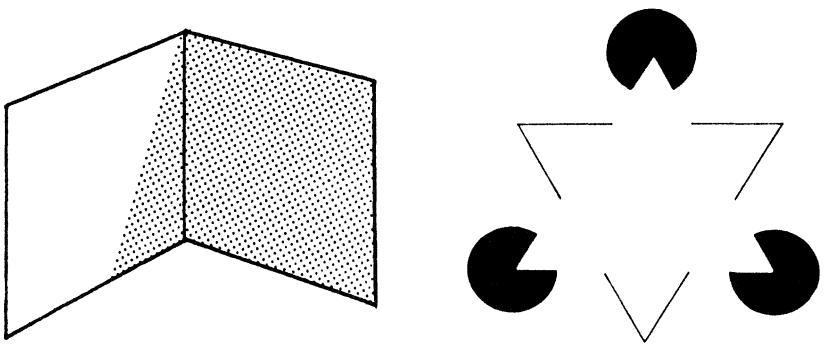

(a) (b)

Figure 1

r/t orientation, or scale, or perspective, or figure/ground, or any of a variety of other dimensions. But in all of these cases one learns very quickly to make the figure flip back and forth at will between the two or more alternatives, by changing one's assumptions about the nature of the object or about the conditions of viewing. At least some aspects of visual processing, evidently, are quite easily controlled by the higher cognitive centers.

One such reversible illusion, which may be unfamiliar to the reader, is striking in that it extends even to changes in perceived color. Take a monochromatic birthday card or similar folded rectangle. Place it upright and oriented to the light so that one of the inside faces is in a very slight shadow relative to the other inside face. (Fig. 1a illustrates the relevant configuration, but only a real card will support the illusion.) Despite this slight shadow, the two faces of the card will be perceived as having the same objective color. Now, closing one eye to defeat stereoscopic orientation cues, treat the object as a Necker Cube and deliberately invert its orientation—in thought—so that the middle fold appears closer to you than the two outside edges. This will produce an obvious distortion in the perceived shape of the card: it will no longer look like a folded rectangle. And it will also produce a change in the perceived color of the shadowed and unshadowed areas of the card. In its original appearance, the slight contrast in luminance is suppressed by the visual system as a mere shadow effect. But in the card's inverted configuration, the slight luminance contrast is no longer consistent with a shadow hypothesis, and the contrast between the two areas is robustly interpreted as a sharp difference in their intrinsic colors. (I owe this example to Richard Gregory.)

Illusory contours provide a similar, but contrasting, example. The white background in Fig. 1b is of course entirely uniform. But most of us can

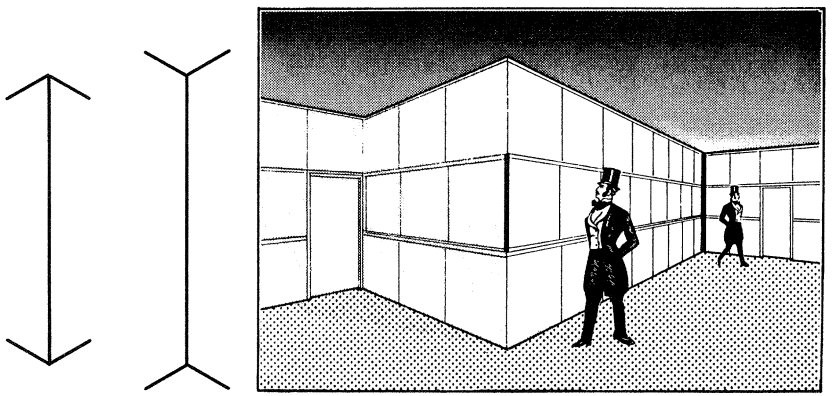

(a) The Müller-Lyer illusion.

(b) The Müller-Lyer illusion in a realistic setting.

Figure 2

see a slightly brighter triangular figure interposed between us and the three black circles; a figure with distinct rectilinear contours marked by a sharp change in luminance, even in the gap between the black circles. Here the eye/brain conjures up luminance differences where in reality there are none. And again, the illusion is penetrable and reversible. Tell yourself that the circles have wedges cut out of them; see the elements of the diagram as six independent objects artfully arranged against a uniform background; center your attention on the two prongs of any V; and the illusory contours disappear.

These assembled examples compile a wide range of elements central to visual perception—contour, contrast, color, orientation, distance, size, shape, figure versus ground—all of which are cognitively penetrable. Collectively, they constitute a strong case against Fodor's claims of impenetrability for our perceptual processing.

But perhaps I am gathering evidence selectively, or aiming it at an exaggerated version of Fodor's view. Perhaps many other elements of perceptual processing, even the dominant share, are *im*penetrable, despite these examples of a contrary cast. What examples does Fodor cite, then, in support of such a claim?

Only one—the Müller-Lyer illusion (Fig. 2a)—though the class he has in mind is clear enough (it will include the Ponzo illusion, the Hering illusion, and similarly persistent illusions). The Müller-Lyer, however, is an odd example for Fodor to be using, because the "text-book story" on how it works (a story apparently endorsed by Fodor, p. 33) explains it as the effect of our having *learned,* in judging absolute size, to make automatic corrections for the variation of an object's angular size with

distance (Fig. 2b.). The illusion exists in the first place only because the relevant processing module is the well-trained victim of some substantial prior education—that is, of some penetration by cognitive activity. The Ponzo and the Hering illusions may have a similar origin. Accordingly, they are all of them poor examples on which to base a general claim of impenetrability.

Now I will grant that, its cognitive origins aside, the Müller-Lyer illusion cannot be overridden by any casual, fleeting, 'voluntary' attempt to modify the character of one's visual experience. By itself, however, this means relatively little, for the issue is not whether visual processing is in general very *easily* or *quickly* penetrated by novel or contrary information: the issue is whether in general it is penetrable at all, where the acceptable means of penetration can include long regimes of determined training, practice, or conditioning. If the Müller-Lyer illusion is an incidental consequence of a long period of perceptual training on certain typical kinds of perceptual problems, then presumably a long period of training in an environment of a quite different perceptual character would produce a subject free from that particular illusion. Fodor, it seems to me, is in no position to insist otherwise, especially given examples of the following kind, which are not speculative, but real.

Recall the effects of chronically worn 'inverting lens' on the visual perception of normal humans. Such lenses have the effect of inverting the orientation of all visual information relative to the body's tactile and motor systems. In short, they turn the visual world upside down. (Kottenhoff 1957, provides a useful summary of this research.)

The initial effect is profoundly disorienting, but with little more than a week's practice, subjects adjust to the new perceptual regime. The subjects are not confined to a chair or bed for the duration of the experiment, but are forced by practical necessity to continue to interact with familiar objects and to engage in the normal forms of motor behavior. The result is that the subjects slowly manage to recoordinate their vision with the rest of their sensory and motor systems, and the illusion of the world's being upside down is said to fade away, all on a time scale of roughly a week.

When the lenses are first put on, and the world is made to appear upside down, the subjects are of course quite aware of what the lenses are doing. They may even know how they do it. But the illusion is not banished by the mere possession of this information. It would clearly be wrong, however, to draw from this any conclusion about the impenetrability of our visual processors. A few weeks of steady practice and experience penetrates them quite nicely. And the degree to which that penetration is successful is further revealed when the lenses are finally removed: for a short time thereafter, the subjects suffer a disorientation illusion very much like

that encountered when the lenses were first put on. Their visual processing, reconfigured by training to compensate for the lenses, continues to 'compensate' after the lenses are gone.

In similar experiments on animals, training produces a reversal in the character of what one might have presumed to be endogenously specified reflexes, such as the vestibular-ocular reflex, which directs one's eyes, when fixated on a target, to move an appropriate amount to the left or right in order to compensate for head movements in the opposite direction. Here the brain seems literally to rewire the relevant neural mechanism under the pressures imposed by left-right inverting lenses (Gonshor and Jones 1976).

Cases like these are important, for they reflect the plasticity of some very deep 'assumptions' implicit in visual processing, such as the specific orientation of the visual world relative to one's other sense modalities and to one's motor systems. If assumptions as deep as these can be reshaped in a week or two, then our perception begins to look very plastic and very penetrable indeed.

I expect Fodor to object, however, that examples such as these, dramatic though they may be, are not cases of the *cognitive* penetration of our peripheral modules. These perceptual changes are wrought not by the simple acquisition of certain beliefs, nor by reflecting on them in the relevant perceptual circumstances. Rather, they are wrought by some form of training, practice, or conditioning, often lengthy.

One way to turn this objection aside is to attack the integrity of the highly questionable dichotomy between 'cognitive penetrability' and other forms of penetrability (see the commentaries on Pylyshyn 1980). But I shall not pursue this path here. There is a simpler and more direct response. *Who ever claimed that the character of a scientist's perception is changed simply and directly by his embracing a novel belief?* None of the theorists cited in Fodor's paper have defended such an unrealistic view. And all of us have, at some point or other, emphasized the importance of long familiarity with the novel idiom, of repeated practical applications of its principles, and of socialization within a like-minded group of researchers.

Kuhn is quite explicit (1962, ch. V, X) that the enveloping paradigm that shapes the scientist's perception is not constituted solely by a set of explicit laws, but by an entire disciplinary matrix that includes standard ways of applying and using the resources of the paradigm, skills acquired during a long apprenticeship. And my own discussion of the plasticity of perception (1979, ch. 2) has the relevant community learning their nonstandard observational vocabulary from birth, in an ongoing *practical* setting where no other idiom is even contemplated.

I confess to having used one example where a temporary shift in per-

ception can be made fairly swiftly—the example of reperceiving the organization of the solar system in a heliocentric rather than a geocentric fashion (1979, pp. 30–34). This case is rather closer to the Necker Cube in character than to the case of the inverting lenses. But even here it was emphasized that simply having the relevant Copernican beliefs is not enough: one must learn how to see the changing heavens as an unfolding instance of the Copernican organization, as viewed from our peculiar perspective within it. Having the relevant beliefs is one thing—we are all of us Copernicans, after all. Reshaping one's perception is quite another.

The point is a general one. A physics student does not come to see the motions of common objects in a new way simply by memorizing Newton's three laws. Most freshman physics students do memorize those laws, but relatively few have their perceptions much altered. The few who do are distinguished by having *practiced* the skills of applying those laws in a wide variety of circumstances. They do come to perceive a common pattern in the behavior of moving bodies that was hitherto invisible to them, but memorizing the laws was only the first step in a fairly lengthy process. There are sudden flashes of insight, to be sure, as when one first grasps how the pattern is instantiated in some typical case. But on the whole, the process of reshaping one's perception takes time, and it requires more than the mere adoption of a belief or three.

To summarize these points: if Fodor is attacking the view that perceptual processing always (or even usually) responds directly and immediately to changes in one's theoretical commitments, then he is attacking a straw man. This is not a view that anyone has defended. On the other hand, if Fodor is denying that perceptual processing is plastic in the face of more comprehensive and protracted kinds of pressures—such as the forced practical *use* of some novel perspective—then the empirical facts are against him. For by these means, even very basic aspects of visual processing can be overturned and reconfigured, as we saw with the visual inversion experiments.

Some degree of "diachronic" penetrability is grudgingly conceded by Fodor (p. 39), since the alternative is to hold that *all* of our adult perceptual capacities are *endogenously* specified. We know that they are not, since the development of so-called 'normal' perception itself plainly involves a great deal of learning on the part of the growing infant. Our perceptual, practical, and social environment shapes our perceptual capacities mightily, especially in their early stages of development, and this suggests that different courses of learning would produce interestingly different perceptual capacities. Fodor attempts to play down this concession, however, by suggesting that the range of possible variation in perceptual development might be quite narrow.

Why he thinks this is left unexplained. The claim needs arguing, the

facts suggest otherwise, and one need not turn to academic journals for shining examples. To see the nonstandard perceptual capacities that our native modalities can acquire, think of the following. In recent centuries, most humans have learned to perceive speech not just auditorally but visually: we have learned to read. Some have learned to perceive speech by touch: they read Braille. And some of us have learned not just to hear music, but to see it: we have learned to sightread musical notation. Now, neither the eyes nor the fingers were evolved for the instantaneous perception of those complex structures and organizations originally found in auditory phenomena, but their acquired mastery here illustrates the highly sophisticated and decidedly supernormal capacities that learning can produce in them. And if these capacities, why not others? Diachronic penetration, I assert, is not only possible and actual, it is commonplace.

Finally, there is neurophysiological evidence that suggests the systematic penetrability of the peripheral modules by the higher cognitive centers. Cell-staining techniques have allowed us to trace out a gross 'wiring diagram' for many parts of the brain. When introduced into a neuronal body, certain chemical stains—notably, horseradish peroxidase—are transported down the entire length of its long fiberlike axon. This marks the axons visually, and their journey through successive sections or slices of the brain can then be followed with an optical microscope. In the case of vision, for example, the dominant nervous pathway starts at the retina, proceeds via the optic nerve to the Lateral Geniculate Nucleus (LGN), and stepwise from there by other pathways to the primary visual cortex, to the secondary visual cortex, and from there to a variety of other areas even higher in the processing hierarchy.

But these 'ascending' pathways are almost invariably matched by 'descending' pathways that lead us stepwise back through the intermediate brain areas. They may even lead back to the earliest processing systems at the retina. The descending projections from the visual cortex back to the LGN, for example, are even greater in number than those in the ascending direction. And though the claim is not well established, there is some evidence that fully 10 percent of the axonal fibers in the human optic nerve are descending projections from the LGN back out to the retinal surface itself, the very first transducer in the processing hierarchy (Wolter 1965; Wolter and Lund 1968; Sacks and Lindenberg 1969).

There are similar chains of descending pathways, from the various areas topmost in the information-processing hierarchy, down through all of the intermediate processing stages and all the way out to the periphery, for all of the other sensory modalities as well. This organizational pattern is typical in mammals, and also in birds (Livingston, 1978, pp. 45–49). Prima facie, the function of these descending pathways is 'centrifugal control'. They allow for the modulation of lower level neural activity as

a function of the demands sent down from levels higher in the cognitive hierarchy. Experimentation on their functional significance is so far limited, but lesions confined to the descending optic nerve pathways (LGN to retina) are known to cause perceptual deficits in birds, even though the descending fibers in their case constitute only 1 percent of the optic nerve total. Lesioned birds are less able than intact birds to distinguish edible seeds from other minute objects in dim light (Rogers and Miles 1972).

If such descending pathways were always sharply confined close to the sensory periphery, or if they were to be found scattered only here and there in the information-processing hierarchy, then we might have some realistic hope of dismissing any backward loop as an element of what is still an 'encapsulated module' from a functional point of view. But descending pathways are the rule in the processing hierarchy of the brain, not the exception. They appear to connect the upper levels in the hierarchy to most and perhaps to all of the lower ones, in each and every one of the sensory modalities. In sum, the wiring of the brain relative to its sensory periphery certainly does not suggest the encapsulation and isolation of perceptual processing. As with the psychological data discussed earlier, it strongly suggests exactly the opposite arrangement.

1.3. Is the Encapsulation Thesis Relevant? Before concluding this section on the etiology of perceptual judgments, I wish to address a further and vitally important point. Let us suppose, for the sake of argument, that perceptual processing is entirely rigid and impenetrable up to the contents of one's visual (auditory, tactile, etc.) manifold; rigid, that is, up to the character of one's sensations. Even if, as now seems very unlikely, visual processing is thus rigid, the outputs of that system are still capable of driving in turn an enormous variety of quite different *conceptual* frameworks. The point here is that sensations themselves are not yet truth-valuable or semantically-contentful states: they are still a stage, though perhaps a late stage, in the processing that leads to specific perceptual judgments or beliefs. Now (and this is a point that I have made explicitly before in 1979, pp. 38–39), however rigid that prior processing might be, there are indefinitely many different possible mappings from the domain of sensations to the domain of propositions (judgments, beliefs), and which of these many mappings comes to characterize your own perceptual activity is a function of which of the indefinitely many conceptual frameworks you have learned as the framework of spontaneous response to the contents of your sensory manifold.

Accordingly, the plasticity of perceptual judgment defended by me in earlier writings does not require that we 'penetrate' the peripheral perceptual modules in any case. We need only *connect* the outputs of those

modules to whatever system of conceptual activity governs our discursive thinking. Further, if two people have learned radically different frameworks, then they will have a severe communication problem despite the rigidity of their peripheral processing. Thus my earlier observation (section 1.1) that encapsulation fails to solve the problem of incommensurability.

That our *conceptual* system is plastic I regard as obvious. That we can successfully connect different conceptual systems to one and the same sensory system I regard as only slightly less obvious. The doctrine of the plasticity of human perceptual judgment requires no premises beyond these.

A few words, then, in support of the less obvious of the two premises. Consider the conceptual framework used for describing pitch in musical theory. It begins with the chromatic scale: C, C♯/D♭, D, D♯/E♭, E/F♭, F, F♯/G♭, G, G♯/A♭, etc. This is not just a list of names. The sequence has a periodic character (octaves), an absolute position in auditory space, and a well-defined metric of various intervals. In fact, the chromatic scale and its various properties form the foundation of musical theory. Clearly, however, this conceptual framework is not innate to our auditory processing, nor is it a part of ordinary language. But people are regularly trained to use it in auditory perception. In time, the better students master what we call a sense of absolute pitch ("That's a middle C, . . . and that's the A♭ above middle C").

More intricately yet, there is the domain of musical chords, and of harmonious sequences of chords. Chords are structured sets of simultaneously sounded notes, sets that fall into an organized matrix of different types (majors, minors, sevenths, ninths, diminisheds, augmenteds, etc.). These also can be directly recognized, by ear, by one suitably practiced in the relevant theory and vocabulary. Such a person perceives, in any composition whether great or mundane, a structure, development, and rationale that is lost on the untrained ear.

We are contemplating a musical example not because it is the only empirical example one can cite, but because it is an unproblematic example. Everyone knows that the 'ear' can be 'trained', as we say, to sustain these remarkable and nonstandard perceptual capabilities. But the example of trained musical perception is a straightforward existence proof for the possibility of theoretically-transformed perception in general. What wants appreciating is that this example is repeatable in a great many other domains. What is required is learning the relevant theory, and extended practice in using it.

We may begin to see some of the endless possibilities by noting that one can just as easily learn to recognize sounds under their dominant *frequency* descriptions as under their music-theoretic descriptions ("That is an oscillation of 262 cycles per second . . . and that's one of 415 cycles

per second"). Equally possible, one can learn to recognize them under their *wavelength* descriptions ("That has a wavelength of 1.19 meters, . . . and that has a wavelength of .75 meters"). The payoff is that one's spontaneous perceptual judgments then put one in a position to anticipate, manipulate, and exploit the details of such auditory phenomena as interference effects, standing waves, Doppler shifts, intensities, and so forth. One requires only a facility with a few elementary laws of wave-propagation.[1]

I conclude this section with an instructive fable. Consider the imaginary community, discussed at length in my (1979), whose members all have the unquestioned assumption that physical objects contain an observable fluid substance, called caloric, which is confined in common objects under a variety of different pressures. According to everyone's 'common sense' convictions, caloric is produced or released in great quantities by fires and by friction; it always flows from high-pressure bodies to low-pressure bodies; a sufficiently high pressure causes the boiling of water; a sufficiently low pressure causes water to freeze; and so forth.

Most importantly, all members of this community regard the pressure of caloric as an observable feature of the world: where you would have, upon touching a simmering kettle, the spontaneous perceptual belief *that this kettle is hot,* they have the spontaneous perceptual belief *that this kettle has a high caloric fluid pressure*. In sum, we have here a community using the conceptual framework of early classical thermodynamics as a common sense *observation* framework for that same range of phenomena commonly addressed by us with 'folk thermodynamics'—the familiar framework of hot and cold. The virtues of this caloric framework, and its persistence in the face of criticism, will not be repeated here, but they are considerable.

I cite this example because the people of this community are making spontaneous 'observational' judgments that are obviously laden with theory. Moreover, the theory at issue is known by us to be false, and so the prejudicial character of all of their perceptual beliefs in this area is made even more dramatic. Consider now a philosopher of this society—Jerry Caloric—who argues as follows.

[1]For example, recall the sound made by a moving car as heard by the pedestrian it passes: *ZEEEEEEYowwwwwww*. Suppose you can hear that the dominant frequency of its approaching hum (*ZEEEEEE*) is 262 hertz (a middle-C), and that its receding hum (*Yowwwwww*) is 220 hertz (a lower A). In such a case you may safely infer that its unshifted or intrinsic frequency must be roughly halfway between the heard extremes, or about 241 hertz. Since the heard frequency is Doppler-shifted from this value by about 21 hertz, which is about 9 percent of the intrinsic frequency (241 hertz), then the velocity of the car must also be about 9 percent of the velocity of sound (740 miles per hour), or about 65 miles per hour. So, if you learn to recognize sounds under their frequency descriptions, then the velocities of unseen objects are often but a quick inference away. Examples like this can be multiplied indefinitely.

"Consider the illusion produced when one's left hand is allowed to rest in a bucket of water at high caloric pressure, and one's right hand in a bucket at low caloric pressure, and then both hands are immersed in a bucket at an intermediate caloric pressure. To judge with the left hand, this water has a low caloric pressure; to judge with the right, it has a high caloric pressure.

"Of course, both we and Granny know how this illusion is produced. The nerves of the two hands become differently fatigued by the extreme caloric pressures in the first two buckets, and thus each gives a different and false response to the intermediate pressure of the third bucket. But notice that possession of this information does absolutely nothing to dispel the illusion.

"We may conclude, therefore, that our peripheral modules are cognitively impenetrable. Accordingly, our perceptual judgments about the *caloric fluid pressures* of common objects are in an important sense theory neutral. The theories we embrace have no effect on caloric perception, and all humans with normal perceptual systems will thus perceive the world in exactly this same way."

I have here recreated the form of Fodor's argument in a setting where the conclusion is clearly false. The point is to highlight some of the ways his argument fails in its original setting. The first lesson is that Fodor, like the philosopher just quoted, fails to appreciate the highly systematic and speculative character of his own observational idioms, a character they will have quite independently of any rigidity in our peripheral modules. And the second lesson is that Fodor never takes seriously the possibility that, even given the rigidity of perceptual processing up to the character of our sensations, one can still train oneself to use, in spontaneous 'observational' mode, conceptual frameworks radically different from those we learned at mother's knee.

2. The Semantics of Observation Predicates. Implicit in the preceding remarks is the view that the meaning of an observation term derives not primarily, nor even perhaps at all, from the typical etiology of its observational application, but rather from the network of general beliefs and assumptions in which it is embedded. Because the contents of such embedding networks can vary substantially, so also can the meaning of our observation terms. Fodor correctly identifies this approach to meaning as a major element in my argument for the theory-laden character of perceptual judgment (I call it "the conceptual role theory of meaning" or "the network approach"; he calls it "meaning holism"). The argument is simple and quickly stated.

(1) Any judgment consists in the application of *concepts* (for example, a is F).

(2) Any concept is a node in a *network* of concepts whose connecting threads are sentences, and its meaning or semantic identity is determined by its peculiar place in that network. (This, in stick-figure form, is the theory of meaning referred to.)

(3) Any network of concepts is a *theory,* minimally, a theory as to some of the classes into which nature divides herself, and some of the relations that hold between them.

∴ (4) Any judgment presupposes a theory.

∴ (5) Any *observation* judgment presupposes a theory.

The theory-ladenness of observation terms thus emerges as a consequence, not of their having some special and regrettable disease, but simply as a consequence of their being meaningful terms at all.

2.1. Objections to the Network Approach: Fodor's Reductio. The defect Fodor finds in an unqualified network approach to meaning is that it allows too much leeway in what an observation sentence might mean.

> 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, in material mode, that *anything might be observed* depending upon theoretical context. (1984, p. 28)

This tracing of presumed equivalences overextends itself. I do not hold that, given normal human senses, *anything* might be observed. You cannot observe what does not exist, and you cannot observe (without instrumental help) what is beyond any physical detection by your native senses. On the other hand, I do assert that almost any predicate could function as the vehicle of spontaneous perceptual judgment for someone trained to conceive of things in the relevant way. But whether his 'observation' judgments constitute genuine cases of veridical perception will be a function of whether (a) the feature he takes himself to be observing really exists, and (b) his sensory system has some reliable discriminatory response to the occurrence of that feature. Failing either of these conditions, his 'observation' judgments will be systematically mistaken, as in the case of the Friends of Caloric. But while mistaken, those judgments may still be highly successful, both from a practical and from a theoretical point of view.

In sum, my position entails that we can observe many features of the world quite different from the features we are used to observing, and that we might not really be observing some of the features that we think we are. But it does not entail that we can observe everything.

2.2. Belief Networks versus Causal Connections. The preceding is not Fodor's main worry, however, and what looked like an attempt at a *reductio* may be just a rhetorical flourish. What he really seems to object to is the idea that the meaning of observation terms might have nothing to do with the objective features of the world that typically elicit or *cause* their spontaneous use. And he cites the possibility that at least some of an observation term's semantic properties might be determined *non*holistically, perhaps by the causal connections just alluded to. He then concludes, "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" (p. 30).

This is not good enough, for two reasons. We do not require Fodor's concession that *all* of the semantic properties of sentences/beliefs are determined by their theoretical context. So long as *some* of the semantic properties of any observation sentence are inevitably determined in that fashion, such sentences will still be stuck with a significant burden of prejudicial theory. To achieve a truly theory-neutral foundation for knowledge, Fodor needs a class of sentences, or terms, *none* of whose semantic properties is dependent on theory.

Second, there are decisive reasons in support of the claim that at least some of any observation term's semantic properties must be determined by the network of beliefs that embeds it. Consider the following argument.

If a term "*F*" is to be a meaningful observation term, then its predication in "*Fa*" must have some material *consequences*: it must imply some further sentences, it must be incompatible with some others, and so forth. The sentence "*Fa*" will clearly have this property if it is asserted in a context where general sentences such as "(x) $(Fx \supset Gx)$", "(x) $((Fx \ \& \ Hx) \supset \sim Kx)$", and so forth, are already assumed. "*Fa*" will then imply "*Ga*", be incompatible with "$(Ha \ \& \ Ka)$", and so forth.

But if "*F*" figures in no such background beliefs or assumptions whatsoever, then "*Fa*" will be entirely without consequence or significance for anything. It will have no bridges to link its assertion or denial with the assertion or denial of any other sentence. It will be a wheel that turns nothing, a coin weightless in every balance, an assertion empty by any measure. Less figuratively, its assertion will be *computationally inert*. It will be without computational significance for the very cognitive system that asserts it.

Meaningful observation terms, therefore, will always be embedded within some set of assumptions. And since there is no analytic/synthetic distinction, those assumptions will always be speculative and corrigible. Meaningful observation terms, we seem bound to conclude, will always be laden with theory.

It will of course be pointed out that, even in the absence of any background assumptions, "*Fa*" will have a host of purely *formal* consequences, such as "$P \vee \sim P$", "$Q \supset Fa$", and so on. But the pattern of these trivial consequences is exactly the *same* for "*Fa*" as it is for any other putative observation sentence—"*Ga*", "*Ha*", etc. These consequences thus cannot serve to bestow any distinct significance on "*Fa*". It is the *material* consequences of "*Fa*" (that is, the ones that flow from substantive or nonformal background assumptions) that do that. Which is another way of stating the central claim of the network theory: what determines the meaning of any term is the peculiar cluster of beliefs in which the term figures, and the peculiar pattern of inferences they make possible.

Despite the clear inevitability of an ideological component in the meaning of any observation term, it remains possible that it is not the only component. Some have argued for a causal component (for example, Putnam 1975, Dretske 1981, Sayre 1986, Fodor 1987). I have criticized causal accounts of meaning elsewhere (Churchland 1983, 1986), and remain inclined towards the idea that causal connections have nothing to do with meaning. But I shall not attempt to argue this stronger claim here. My aim in this section is only to reestablish the strong and independent presumption that any observation term, to the extent that it is meaningful at all, must be embedded in a network of corrigible assumptions.

2.3. Sensational Plasticity versus Conceptual Plasticity. One possible way to defend Fodor would be to concede the theory-dependent character of our observational concepts and judgments, and try to insist on no more than the theory-independent character of our sensations. Fodor himself seems to be sketching a position of this sort late in his paper, when he urges the rigidity of "the look of things" versus the plasticity of "how things are judged to be" (p. 40).

But this defense will not take us any distance at all. For one thing, if all Fodor wishes to insist on is uniformity in the character of our sensations through changes in our doxastic commitments, then his argument is largely an *ignoratio*. It fails to address the major epistemological tradition at issue, whose central theme has always been the theory-laden character, not of our sensations, but of our observational concepts and observational judgments.

And there is a very good reason for the centrality of that theme. Thinkers in the tradition at issue (Popper, Feyerabend, Hanson, etc.) have been primarily concerned with the refutation or corroboration of theories. But sensations themselves neither confirm nor refute any theory. Sensations belong to the wrong logical space: it is only an observation *judgment* or *belief* or *report* that can be logically consistent or inconsistent with any

theory (Popper 1959). Thus the chronic concern, throughout the positivist and post-positivist periods, with the possibility of a theory-neutral observation *vocabulary*. Whether sensations themselves might be infected or modified by theory was rarely, if ever, an issue.

My own 1979 position, to cite one target of Fodor's, simply assumes the generally constant character of our sensory responses to the environment. The plasticity that excited me there was confined to the *conceptual* frameworks within which we make our judgmental responses to the passing contents of our sensory manifold. Accordingly, if rigidity in the character of our sensations is all Fodor is concerned to defend, then I do not understand his objection to and dismissal of (p. 28–29) the alternative perceptual possibilities sketched in my 1979 (p. 30). For that sketch makes no assumptions about the plasticity of our sensations. It is *conceptual* plasticity that is there at issue.

To be sure, sensational plasticity would constitute an *additional* argument for the plasticity of perception. At least one author has cautiously advanced a claim of this kind (Kuhn 1962, p. 120–121). And I, for another, am now willing to defend it vigorously (recall the examples in Fig. 1). So there is a genuine point to attacking it, as Fodor does. But it is wrong to represent or to regard this attack, successful or otherwise, as aimed at the principal arguments in favor of theory ladenness. Those arguments have typically been based on other grounds entirely: on the plasticity of our conceptual responses to sensory activity.

3. Conclusion. I will spare the reader a reprise of the various conclusions already reached. Instead, let me try to evoke a general picture of the situation. The central issue of this paper is not an argument about the obscure etiology of a certain class of beliefs, or the arcane semantics of a certain class of terms. The real disagreement is about the fundamental character of the human epistemic situation, and the long-term possibilities for the evolution of the human spirit.

Our epistemic situation, I assert, is one in which even the humblest judgment or assertion is always a speculative leap, not just in its assertion over its denial, but also in the background conceptual framework in which that judgment is constituted, in preference to the infinity of other conceptual frameworks that one might have used instead. In the case of perceptual judgments, what the senses do is cause the perceiver to activate some specific representation from the antecedent system of possible representations—that is, from the conceptual framework—that has been brought to the perceptual situation by the perceiver. A perceptual judgment, therefore, can be no better, though it can be worse, than the broad system of representation in which it is constituted.

This means that perceptual judgments are evaluable at two distinguish-

able levels. The first concerns the propriety of the judgment as evaluated by the local standards of the framework that embeds it (Was the observer in a position to make it? Was the observation made carelessly? Is it inconsistent with information already in hand?). These correspond to what Carnap (1956) has called "internal questions". The second level concerns the adequacy of the embedding framework overall, as a system adequate to represent the range of nomologically possible configurations that the objective world might assume. These correspond roughly to what Carnap has called "external questions" (Do the categories of my framework capture the objective divisions in Reality? Do the basic generalizations of my framework express genuine laws of Nature?).

'External questions' are rather daunting. Local standards of evaluation are both inapplicable and question begging, and global standards are vague and elusive. Good positivist that he was, Carnap reacted by denying that such questions are factual in character, claiming that the decision to use a given conceptual framework is ultimately just a practical question, to be decided on pragmatic grounds. But Carnap was mistaken in seeing a fundamental difference between the two kinds of questions, and between two kinds of grounds for acceptance. So-called 'external questions' are just large-scale theoretical questions, to be decided on empirical and systematic grounds like anything else. And so-called 'pragmatic' considerations attend epistemic decisions at every level of inquiry, even the most humble and mundane.

In fact, 'external questions' are confronted by humans and dealt with on a daily basis, by scientists inventing and evaluating new frameworks for understanding this or that domain, and also by infants and children, who must evolve a conceptual framework adequate to conduct a life in the *lebenswelt* of concurrent human society. This means that our conceptual frameworks can and regularly do undergo *change,* both within the lifetime of an individual, and in society as a whole, over historical periods. To use a Hegelian figure, the journey of the human spirit is essentially the story of our evolving conception of the world, and of our own place within it. Our eyes are little different from a Baboon's or a Chimpanzee's, but our perceptual knowledge is profoundly superior to theirs. Our motor systems are little different from those of any other primate, but our practical capabilities and intentional actions encompass universes quite closed to them. The main difference lies in the dramatically superior conceptual frameworks we have evolved—epigenetically, and not without misadventure—over the course of the last 500,000 years.

If we have come this far, must the journey end here? Manifestly not. The long awakening is potentially endless. The human spirit will continue its breathtaking adventure of self-reconstruction, and its perceptual and motor capacities will continue to develop as an integral part of its self-

reconstruction. But only if we try hard to see new opportunities, and only if we work hard at leaving old frameworks behind.

REFERENCES

Bruner, J. (1973), "On Perceptual Readiness", in J. Anglin (ed.), *Beyond the Information Given*. New York: W.W. Norton & Co., pp. 7–42.

Carnap, R. (1956), "Empiricism, Semantics, and Ontology", in R. Carnap, *Meaning and Necessity*, second edition. Chicago: Chicago University Press, pp. 205–221.

Churchland, P. M. (1975), "Two Grades of Evidential Bias", *Philosophy of Science 42*: 250–259.

———. (1979), *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

———. (1986), "Semantic Content: In Defense of a Network Approach", *The Behavioral and Brain Sciences 9*: 139–140.

Churchland, P. M. and Churchland, P. S. (1983), "Information: Semantic and Information-Theoretic", *The Behavioral and Brain Sciences 6*: 67–86.

Dretske, F. (1981), *Knowledge and the Flow of Information*. Cambridge, Mass.: The MIT Press.

Fodor, J. (1984), "Observation Reconsidered", *Philosophy of Science 51*: 23–43.

———. (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass.: The MIT Press.

Gonshor, A., and Jones, G. M. (1976), "Extreme Vestibulo-Ocular Adaptation Induced by Prolonged Optical Reversal of Vision", *Journal of Physiology 256*: 381–414.

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

———. (1974), *Concepts and Mechanisms of Perception*. New York: Charles Scribner's Sons.

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

Kottenhoff, H. (1957), "Situational and Personal Influences on Space Perception With Experimental Spectacles", *Acta Psychologica 13*: 79–97.

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

Livingston, R. B. (1978), *Sensory Processing, Perception, and Behavior*. New York: Raven Press.

Popper, K. (1959), *Logic of Scientific Discovery*. New York: Harper & Row.

Putnam, H. (1975), "The Meaning of 'Meaning'", in K. Gunderson (ed.), *Language, Mind and Knowledge, Minnesota Studies in the Philosophy of Science, Vol. 7*. Minneapolis: University of Minnesota Press.

———. (1981), *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Pylyshyn, Z. (1980), "Computation and Cognition: Issues in the Foundation of Cognitive Science", *Behavioral and Brain Sciences 3*: 111–134.

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

Rogers, L. J., and Miles, F. A. (1972), "Centrifugal Control of Avian Retina: Effects of Lesions of the Isthmo-Optic Nucleus on Visual Behavior", *Brain Research 48*: 147–156.

Sacks, J. G., and Lindenberg, R. (1969), "Efferent Nerve Fibers in the Anterior Visual Pathways in Bilateral Congenical Cyctic Eyeballs", *American Journal of Opthalmology 68*: 691–695.

Sayre, K. M. (1986), "Intentionality and Information Processing: An Alternative Model for Cognitive Science", *The Behavioral and Brain Sciences 9*: 121–138.

Wolter, J. R. (1965), "The Centrifugal Nerves in Human Optic Tract, Chiasm, Optic Nerve, and Retina", *Transactions of the American Opthalmological Society 63*: 678–707.

Wolter, J. R., and Lund, O. E. (1968), "Reaction of Centrifugal Nerves in the Human Retina", *American Journal of Opthalmology 66*: 221–232.