
A Reply to Churchland's "Perceptual Plasticity and Theoretical Neutrality"

Author(s): Jerry A. Fodor

Source: *Philosophy of Science*, Vol. 55, No. 2 (Jun., 1988), pp. 188-198

Published by: The University of Chicago Press on behalf of the Philosophy of Science Association

Stable URL: <https://www.jstor.org/stable/187957>

Accessed: 22-10-2019 16:08 UTC

REFERENCES

Linked references are available on JSTOR for this article:

https://www.jstor.org/stable/187957?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Philosophy of Science Association, The University of Chicago Press are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*

A REPLY TO CHURCHLAND'S "PERCEPTUAL PLASTICITY AND THEORETICAL NEUTRALITY"*

JERRY A. FODOR

*Department of Philosophy
The Graduate Center, City University of New York*

Churchland's paper "Perceptual Plasticity and Theoretical Neutrality" offers empirical, semantical and epistemological arguments intended to show that the cognitive impenetrability of perception "does not establish a theory-neutral foundation for knowledge" and that the psychological account of perceptual encapsulation that I set forth in *The Modularity of Mind* "[is] almost certainly false". The present paper considers these arguments in detail and dismisses them.

I have it in mind one of these days to write a paper called "Modularity and Objectivity" (or maybe "Objectivity and Modularity"). This, however, isn't it. What I propose to do in this paper is argue a *very* narrow case. Churchland offers a batch of considerations intended to convince us that the cognitive impenetrability of perception "does not establish a theory-neutral foundation for knowledge" and that my empirical "views on impenetrability are almost certainly false". I propose to go through these arguments and show, in some detail, that they are no good; that is, that they are no good *whether or not their conclusions are true*.

Churchland's paper is mostly concerned with three topics: 1. What are the epistemological implications of perceptual encapsulation (assuming, for the moment, that perceptual processes are indeed encapsulated)? 2. Is the encapsulation thesis true? 3. Some semantical considerations that are supposed to show that the meaning of observation terms must be theory dependent even if the perceptual processes involved in observing things are encapsulated and theory neutral. I propose to discuss Churchland's arguments under these heads, but with a spare category for miscellanea.

1. The Epistemological Implications of Encapsulation

Churchland: "Let us suppose . . . that our perceptual modules. . . . embody a systematic set of . . . assumptions about the world, whose influence on perceptual processing is unaffected by any additional or contrary information. . . . this may be a recipe for a certain limited *consensus* among human perceivers, but it is hardly a recipe for theoretical

*Received April 1987; revised May 1987.

Philosophy of Science, 55 (1988) pp. 188–198.
Copyright © 1988 by the Philosophy of Science Association.

neutrality. . . what we have is a universal dogmatism, not an innocent Eden of objectivity. . . Encapsulation does nothing to ensure the truth of our perceptual beliefs . . ." (pp. 169–170).

Reply: Nobody was offering innocence or a guaranty of truth. The question at issue is: What are the psychological conditions under which differences among the theories that observers hold are *not* impediments to perceptual consensus among the observers. Cognitive encapsulation seems to be an empirically necessary condition for this, and one that is (contrary to New Look psychologizing) apparently satisfied.

However, if you consider the sort of background information that penetrates perception (according to modularity theory), it turns out that perception *is* neutral, de facto, *with respect to most of the scientific (and, for that matter, practical) disagreements that observation is called upon to resolve*. According to standard versions of modularity theory (including the version I set out in *Modularity of Mind*) perceptual processing has access only to background information about certain pervasive features of the relations between distal layouts and their proximal projections. (Hardly surprising, since it is precisely the relation between proximal and distal stimuli that perceptual processes are required to compute.) Thus, in the case of vision, a good candidate for accessible background is information about the geometrical relations between three-dimensional objects and the two-dimensional images they project onto the surface of the retina. In the linguistic case, a good candidate for accessible background is information about the grammatical structures that inform the type/token relation for the speaker/hearer's dialect.

The point is that, in both cases, reliance on such information constitutes a perceptual bias; and in both cases it makes perception 'inferential' in the required sense. But this bias leaves perception neutral with respect to *almost all* theoretical disputes, so it couldn't ground any *general* argument for the unreliability of observation. Contrary to Churchland, there seems no reason to doubt that this very restricted sort of bias might be compatible with more than enough perceptual neutrality to "secure for us any theory-neutral foundation for knowledge" (p. 171). Indeed, it might leave us with enough theory-neutral observation to allow us *to discover, and correct for, our own perceptual biases*. We might do so by relying upon inferences from theories to the observational confirmation of which our perceptual biases are irrelevant. This sort of boot strapping is complicated to describe but often routinely easy to perform.

By the way, the preceding is not me pulling in my horns after the fact. That the premises to which perceptual inferences can appeal are substantively restricted by the architecture of the mind is the *whole point* of modularity theory.

Churchland: “[The] consensus would last only until the first mutant or alien comes along, to confront us with a different perceptual point of view” (p. 171).

Reply: Churchland apparently wants a naturalistic account of scientific objectivity to supply a guaranty that an arbitrary collection of intelligent organisms (for example, a collection consisting of some *Homo sapiens* and some Martians) would satisfy the empirical conditions for constituting a scientific community. *Of course* there can be no such guaranty. Our dependence upon the reliability of our cognitive faculties—perceptual biases and all—is part of the inductive risk that makes scientific inference non-demonstrative. On the other hand: I once had a book that purported to divide all the possible worries into the Real and the Merely Baroque. Churchland’s worry—that (unspecified) aliens might arrive at a science different from ours in virtue of (unspecified) differences between their perceptual biases and ours—belongs, it seems to me, to the second category.

2. Is the Encapsulation Thesis True? Two preliminary points: first, modularity is an *empirical* thesis, so how it comes out depends largely on what the psychological data prove to be. Second, the epistemologically relevant question is not whether modules are perfectly encapsulated, but whether they are encapsulated enough to permit theory-neutral, observational resolution of scientific disputes. Now read on, *s.v.p.*

2.1. Ambiguity

Churchland: “. . . many illusions . . . [show] . . . that our visual modules are indeed penetrable by higher cognitive assumptions. . . . one learns very quickly to make the [ambiguous] figure flip back and forth at will . . . by changing one’s assumptions about the nature of the object or about the conditions of viewing” (pp. 171–172).

Reply: False. One doesn’t get the duck rabbit (or the Necker Cube) to flip by “changing one’s assumptions”; one does it by (for example), changing one’s fixation point. Believing that it’s a duck doesn’t help you see it as one; wanting to see it as a duck doesn’t help much either. But knowing where to fixate *can* help. Fixate there *and then the flipping is automatic*.

When one becomes sophisticated about the laws that govern the way things look, one can finagle the looks by playing the laws. In the most obvious cases: one squints to make things look sharper; one cups one’s hand behind one’s ear to make them sound louder, etc. It doesn’t begin

to follow that auditory and visual acuity are cognitively penetrable.

Exactly similarly, one learns that one can get the figure to flip by altering one's fixation point (or, for that matter, by just *waiting*; eventually it will flip of its own accord). To confuse this with the penetration of perception by utilities is to make the following mistake:

a: Heart rate is cognitively penetrable! I can choose the rate at which my heart beats.

b: Remarkable; how do you do it?

a: Well, when I want it to beat faster, I touch my toes a hundred times. And when I want it to beat slower, I take a little nap.

b: Oh.

Churchland has some further, rather complicated cases on offer in which the reversal of an ambiguous figure brings other perceptual effects automatically in train (for example, if you see the figure as reversed in depth, its apparent surface illumination is also seen to change). Churchland's conclusion seems to be: So I can see the surface illumination as I choose.

But these examples don't advance the argument; they rest on the same mistake just scouted—only, as it were, at one further remove. What is going on is: (a) there's a choice about how you see the shape-ambiguous figure; and (b) there's a nomic connection between seeing the figure as having a certain shape and seeing it as having a certain surface illumination. So you get to see the illumination you want *by* choosing how you see the shape. (And you get to see the shape you want by, for example, squinting, altering your fixation point, etc.) *It doesn't follow that you can choose how you see the illumination*; all that follows is that there are things you can do to get yourself to see the illumination one way or the other. (Compare the heartbeat case.) A fortiori, it doesn't follow that there are "a wide range of elements central to visual perception . . . all of which are cognitively penetrable" (p. 173). Indeed, so far, we haven't seen *any*.

It may be that you can resolve an ambiguous figure by deciding what to attend to. But (a) which figures are ambiguous is *not* something you can decide; (b) nor can you decide what the terms of the ambiguity are; (c) nor can you decide what further psychological consequences (for example, consequences for apparent illumination) the resolution of the ambiguity will entrain. This all sounds pretty unpenetrated to me.

Attention is, in short, a wild card in an account of observational neutrality; but it may well be that if you fix the perceptual apparatus and you fix the object of attention, then you fix the appearances for all normal observers *even in the case of ambiguous figures*. If this is true it's episte-

mologically interesting since part of arriving at a consensus as to what experiment to perform to choose among rival theories is agreeing about *what part of the experimental environment to attend to*. “It’s where the dial points to that matters, not the color of the numerals”; and so forth.

Final word about Necker Cubes. Even if they showed that the perceptual analysis of structurally ambiguous figures is unencapsulated (which they don’t), that mightn’t matter much for the neutrality of observation at large since, patently, *most stimuli aren’t structurally ambiguous*.

2.2. Synchronic and Diachronic Penetration

Churchland: 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 . . . [for example by] . . . some form of training, practice or conditioning, often lengthy” (pp. 174–175).

Reply: It looks to me as though there are several issues. Let’s see where we are.

It used to be thought that there is lots of evidence for relatively short-term effects of beliefs and utilities on perception; perceptual effects of your expectations about the color/suit correlations of playing cards; perceptual effects of transient peer pressures, etc. This was the evidential stuff of which New Look perceptual theory was made. And it was worrying because insensitivity to local alterations in beliefs and utilities is, in any event, a *necessary* condition for the theory neutrality of observation.

But now it is conceded that there may, after all, be *no* such local effects. It is, perhaps, only “comprehensive and protracted kinds of pressures” (p. 176) to which perceptual processing is plastic. (These might not even be perceptual effects of acquiring *beliefs*; perhaps they’re perceptual effects of *having the experiences* in virtue of which the beliefs are acquired.)

How much would this matter? What degree of *diachronic* encapsulation would be required for the possibility of theory-neutral observational resolution of scientific disputes? Well, surely less than cast-iron insensitivity of perceptual processes to training. Rather, what seems to be required is just enough diachronic encapsulation to allow perceptual consensus to survive the effects of the kinds of differences of learning histories that observers actually exhibit. For example, if training affects perceptual acuity, then that would be a kind of failure of diachronic encapsulation; but it wouldn’t be anything that an epistemologist need worry about since observational consensus doesn’t generally depend on the observers all

having perceptual acuity to the same degree.

Well, what's the evidence? Is there enough diachronic encapsulation for the purposes at hand?

Answer: moot. Naturalized epistemology awaits the empirical findings. Whereas there's a respectable empirical argument to be made for synchronic encapsulation, *nobody* knows what's going on in the diachronic case; the only point that is worth making is that *if* diachronic encapsulation proves to be pervasive, then we will be within hailing distance of a naturalistic account of how theory-neutral observation is possible.

In any event, the point of present concern is that the considerations Churchland raises as militating against diachronic encapsulation cut next to no ice at all. There are a number of these:

2.2.1. Inverting Lenses. It is, at first blush, a shock to modularity theory that people can adapt to such drastic affronts to their perceptual prejudices as the inversion of the retinal image. This really does suggest the sort of perceptual plasticity—the sort of penetration of perception by experience—that modularity theory says shouldn't be there.

That's first blush; second blush is much better. For there are, after all, good ecological reasons why you might expect plasticity of this sort. Namely: organisms *grow*, and as they grow they must recalibrate the perceptual/motor mechanisms that correlate bodily gestures with perceived spatial positions (paradigmatically, in the human case, the mechanisms of hand/eye coordination). That is, what needs to be kept open for recalibration is whatever mechanisms compute the appropriate motor commands for getting to (or pointing to, or grasping) a visible object on the basis of its perceived location. Adaptation to inverted (and otherwise spatially distorting) lenses is plausibly an extreme case of this sort of recalibration. Indeed, there is experimental evidence that this is so. It turns out that smooth adaptation occurs only when the subject is permitted to actively manipulate the environment. In particular, adaptation does *not* occur (much) in organisms that are, for example, passively wheeled around but deprived of perceptual-motor feedback. (See Held and Bossom 1961.)

In short, the subject in an inverting lens experiment has to learn such things as to *grasp* down for what *looks* up and vice versa. And this sort of relearning is likely not different in kind from the corrections that have to be made for alterations in the angular relations between hand, eye and distal object in consequence of growth. So it's plausible that there are specific mechanisms that function to effect the required visual-motor calibrations, and that it's these mechanisms that are engaged in adaptation to inverting lenses. The moral of the inverting lens experiment thus seems to be: you find specific perceptual plasticity pretty much where you'd expect to find it on specific ecological grounds. What Churchland needs

to show—and doesn't—is that you *also* find perceptual plasticity where you *wouldn't* expect it on specific ecological grounds; for example, that you can somehow reshape the perceptual field by learning physics. Churchland offers, however, no examples of this. I strongly suspect that's because there aren't any.

2.2.2. Reading

Churchland: “In recent centuries . . . [we] . . . have learned to perceive speech, not just auditorally but visually: we have learned to read. . . . Now . . . the eyes . . . were [not] evolved for the instantaneous perception of those complex structures . . . originally found in auditory phenomena, but their acquired mastery here illustrates the highly sophisticated and . . . supernormal capacities that learning can produce in them” (p. 177).

Impatient reply: In recent centuries we have learned to perceive automobiles (not just auditorally but visually). Now the eyes were not evolved for the instantaneous perception of those complex structures. So doesn't their acquired mastery illustrate the highly sophisticated and supernormal capacities that learning can produce in perception?

Fiddlesticks. Churchland needs, and doesn't have, an argument that the visual perceptual capacities of people who can read (or, *mutatis mutandis*, people who can automobile spot) differ in any interesting way from the visual perceptual capacities of people who can't. In precisely what respects does he suppose illiterates to be *visually* incapacitated?

The old story is: you read (spot automobiles) by making educated inferences from properties of things that your visual system *was* evolved to detect; shape, form, color, sequence and the like. Churchland offers no evidence that educating the inferences alters the perceptual apparatus.

2.2.3. Neurological Data

Churchland: There are lots of “descending pathways” from higher to perceptual centers. To be sure, “experimentation on their functional significance is so far limited, but . . . ” (p. 178).

Reply: None required. Heaven knows what psychological function “descending pathways” subserve. (Heaven knows what psychological function 98.769 percent of known neuroanatomical structures subserve, for that matter.) One thing *is* clear: if there is no cognitive penetration of perception, then at least “descending pathways” aren't for *that*.

2.2.4. *Perceptual Learning*

Churchland: someone musically sophisticated “perceives, in any composition whether great or mundane, a structure, development, and rationale that is lost on the untrained ear” (p. 179).

Reply: This merely begs the question, which is whether the effects of musical training are, in fact, perceptual. Churchland adds that “one can just as easily learn to recognize sounds under their dominant *frequency* descriptions . . . [or] . . . under their *wavelength* descriptions” (pp. 179–180), but again no argument is provided that someone who has learned this has learned to perceive differently (as opposed to having learned a different way of labeling his perceptions and a different theory about what his perceptions are perceptions *of*; see section 3.1 below).

What Churchland has to show is, first, that *perceptual* capacities are altered by learning musical theory (as opposed to the truism that learning musical theory alters what you know about music;) second, that it’s learning the theory (as opposed to just listening to lots of music) that alters the perception; and third that perception is altered in some different way if you learn not musical theory but acoustics. Churchland doesn’t show any of these things—he doesn’t even bother to *argue* for any of them—and I doubt that any of them are true. (Attempts to make a case for the corresponding phenomena in color perception have not fared well; see the recent experimental literature on the “Whorf hypothesis”.) In any event, you don’t refute modularity theory by the unsupported assertion that it is contrary to the facts.

3. Miscellaneous: Two Digressions

3.1. *The Argument about Caloric*. I am not at all clear how Churchland thinks this argument goes. I paraphrase under correction.

Churchland: Somebody who describes his heat experiences in terms of caloric theory could insist upon the cognitive impenetrability of ‘caloric illusions’; (for example, of the two bucket illusion). With the absurd conclusion that “our perceptual judgments about the *caloric fluid pressures* of common objects are in an important sense theory neutral” (p. 181).

Reply: What on Earth does Churchland suppose that this observation shows? The theory neutrality of perception isn’t about the impact of one’s beliefs upon how one *describes* one’s experiences; it’s about the impact

of one's beliefs upon one's experiences. It is thus perfectly true, and perfectly harmless, that our perceptual judgments about the caloric fluid pressures of common objects are in an important sense theory neutral; that is, they are theory neutral qua perceptual judgments, but not qua judgments about caloric fluid pressures. Thus, if we changed theories, then we would no longer describe the illusion in terms of the apparent caloric pressures in the two buckets; perhaps we'd describe it in terms of the apparent mean molecular kinetic energy (mmke). But, to repeat, the encapsulation thesis isn't that changing a guy's beliefs leaves his *descriptions* of his experiences intact; it's that it leaves *the experiences themselves* intact; in the present case, changing from the caloric story to the mmke story *doesn't make the illusion go away*.

I do not wish to seem to harp on this, but *really!* The 'false' conclusion of which the thought experiment is supposed to be a reductio is that "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" (p. 181). Now, (a) the first conjunct is surely true; since there is no such thing as caloric, there is no such thing as caloric perception. What theories one holds doesn't change that, so the theories we embrace have no effect on caloric perception. And (b) the second conjunct *may* be false, but it's not shown to be by remarking that if you think there is caloric and you don't think there is mmke, then if you have a heat illusion you will describe it as a caloric illusion and you won't describe it as an mmke illusion. It's not only not *shown*; the observation doesn't even *bear*.

If you experience a perceptual phenomenon, and you happen to think it's the sort of perceptual phenomenon that Granny is always experiencing, then you will perhaps describe it as a Granny phenomenon. And if you then happen to stop thinking that it is the sort of phenomenon that Granny is always experiencing, you will then perhaps stop describing it as a Granny phenomenon. These truisms do *not* tend to substantiate the hypothesis that your perceptual phenomena are penetrated by your beliefs about Granny. (Or, for that matter, to substantiate its denial.)

It may be that Churchland has in mind an argument that goes like this: Our theories change the way we *describe* our experiences. But establishing a scientific consensus requires that there be some descriptions of perception that are theory neutral (for example, the dial is pointing to the "7"; the fluid has turned pink; etc.). So even if our *experiences* are theory neutral, that's not enough for theory-neutral observational validation of our theories; not, at least, if observational validation is something that scientific *communities* do.

Reply: The thought experiment about caloric shows that *some* of the ways we describe our experiences change with changes in theory (so does

the thought experiment about Granny); but what Churchland needs is that *all* of the ways we describe our experiences are (in principle) theory sensitive. In effect, he needs to argue that there can be no theory-neutral observation vocabulary even if there is theory-neutral observation. This seems to me, to put it mildly, less than self-evident. In any event, it surely does not follow from the thought experiments. Or from any other argument that Churchland offers, so far as I can tell.

3.2. Digression on Sensations

Churchland: “. . . 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. . . alternative perceptual possibilities . . . [that make] no assumptions about the plasticity of our sensations” (p. 185).

Reply: Churchland constructs a sensation/judgment dilemma, and then proposes that I impale myself on one of the horns. No thanks. There may be some nontruthvaluable (purely sensory) states involved in perception, but they aren't the output states of modules. To a first approximation, the outputs of modules are judgments about how things appear; judgments which are then up for being interpreted and corrected by reference to background beliefs in the course of 'higher' cognitive processing. The idea is that there are two sorts of judgmental processes (perceptual and higher cognitive) one but not the other of which is encapsulated. This idea is *neutral* on the issue whether there is also some *nonjudgmental* process whose encapsulation might follow (perhaps trivially) from its nonjudgmentalness. Modularity theory is neutral on all of this, and so am I.

4. Semantics

Churchland: If you accept a 'conceptual role' story about meaning, then it will probably follow that what theory you hold determines what your observation statements mean.

Reply: So much the worse for conceptual role stories about meaning. So much the worse for use theories in general, for that matter; I wouldn't have one at a discount.

Churchland: You had better accept a conceptual role story about meaning, because “If a term '*F*' is to be a meaningful observation term,

then it's predication in '*Fa*' must have some material *consequences*: it must imply some further sentences. . . . But if '*F*' figures in no . . . background beliefs or assumptions whatsoever, then '*Fa*' will be entirely without consequence or significance for anything. . . . It will be a wheel that turns nothing. . . . Meaningful observation terms, therefore, will always be embedded within some set of assumptions. And since there is no analytic/synthetic distinction, these assumptions will always be speculative and corrigible" (p. 183).

Reply: (a) From the fact that meaningful observation (or other) terms are always embedded in a theory, it does not follow that the theory that a term is embedded in contributes to determining what it means. (b) The observation sentence '*Fa*' is true iff *a* is *F*. So, by assumption, '*Fa*' has a truth condition and is a fortiori significant. It would appear that this is so *whether or not* '*F*' "figures in background beliefs or assumptions", so I'm at a loss to imagine what argument Churchland thinks he has given for a conceptual role theory of meaning. (Of course, Churchland might claim that '*Fa*' *couldn't* have a truth condition *unless* '*F*' figures in background beliefs; but that would be to beg the question and establish conceptual role semantics by fiat.) For discussion of what appears to be a similar bad argument that turns up in Dennett's "Intentional Systems", see my *Psychosemantics*, p. 89.

5. Coda

Churchland: ". . . must the journey end here? . . . 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" (pp. 186–187).

Reply: An endless awakening sounds like not all that much fun, come to think of it: I, for one, am simply unable to self-reconstruct until I've had my morning coffee. Actually, theories come and theories go and people don't really change very much; or so it seems to me. That's probably just as well; if we *become* our theories, how are they to "die in our stead"?

REFERENCES

- Churchland, P. M. (1988), "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor", *Philosophy of Science* 55: 167–187.
 Fodor, J. (1983), *The Modularity of Mind*. Cambridge: The MIT Press.
 ———. (1987), *Psychosemantics*. Cambridge: The MIT Press.
 Held, R. and Bossom, J. (1961), "Neonatal Deprivation and Adult Rearrangement: Complementary Techniques for Analyzing Plastic Sensory-Motor Coordinations", *Journal of Comparative and Physiological Psychology* 56: 872–876.