Broadening the Mind

John Perry
Stanford University

July 3, 1997

[Review of Jerry Fodor, *The Elm and The Expert*. *Philosophy and Phenomenological Research*, 1997.]

The main topic of Jerry Fodor's *The Elm and the Expert*,[1], and the title of the first chapter, is "If Psychological processes are computational, how can psychological laws be intentional?" I focus on the first and second chapters; The first is devoted to setting up the question, the second to answering it.

1 The Problem

The topic is an old one for Fodor— as he says he has been thinking and writing about it for more than twenty years. What seems to have been constant since, say, "Methodological Solipsism," is his devotion to theses that we'll call, following his scheme in the book, (1) and (3):

(1) Psychological explanations are intentional, that is, backed by laws that classify by content.

(3) Psychological laws are implemented by computational processes.

What has changed is Fodor's view about the issue of broad and narrow content. The contents in question are those of linguistic acts and cognitive states. In the case of an assertion or a belief, contents will be or be closely related to truth-conditions: what conditions must the rest of the world fulfill, in order for the assertion or belief to be true? Content is narrow or broad, depending on how much we include in "the rest of the world," that is, what facts we take to be fixed and what we allow to vary. Take the assertion "The author of *The Elm and The Expert* likes to sail." It can be assigned two quite different contents, depending on how one thinks of the description. If we fix the fact that Fodor is the author, then in order for the assertion to be true, he has to like to sail. If we don't take that fact as fixed, then we get that in order for the assertion to

[1] Jerry Fodor, *The Elm and the Expert* (Cambridge, Massachusetts and London, England: MIT-Bradford, 1994), pp. xii, 129.

be true there has to be someone who is both the author and likes to sail. The second content is *narrower* in that it fixes less and allows more to vary.

Seminal work in the philosophy of language and the philosophy of mind in the 1970's argued persuasively that the content we ascribe to linguistic acts (what is said) and beliefs (what is believed) were, by and large, and in various ways, broad.

The work of Kripke, Donnellan, Kaplan and others on the "New Theory of Reference," for example, showed that our ordinary conceptions of what is said and what is believed, in the case of assertions involving names and indexicals and the beliefs they express, are broad in the sense that they involve the individuals named or contextually indicated, rather than descriptions or conceptions of them. That is, the truth-conditions take the facts of reference involving names and indexicals as fixed, rather than things allowed to vary.

The work of Putnam, Burge and others focused on properties rather than individuals. When Elwood and Telwood use the term "water," it is not their ideas or mentalese terms for water (which are the same) that enter into what they say and believe. It is the actual kind of stuff, water or twater, that they are talking about. When Burge's patient believes that he has arthritis in his thigh, he believes something false, and the false thing he believes is a proposition about arthritis. This is broad content because it is the conventions of his particular linguistic group, even if only partly reflected in his own understanding, that determine which kind of inflammation he speaks about and thinks about with the word "arthritis."

I'll call these currents from the 70's "referentialism." Referentialism has coalesced in various ways with the informational tradition in semantics, in ways that Fodor seems to embrace but doesn't trace. By broad content, he means content that involves the properties, relations, and individuals that the ideas are *of* or the words are *about*. Which properties, relations and individuals depends not only on the nature of the ideas or words and their connections to other words and ideas, but on causal relations with their external referents.

In "Methodological Solipsism," Fodor argued that theses (1) and (3) nevertheless showed the importance of a suitably narrow level of content. Since content supervenes on computational processes, it must be narrow, i.e., not depend on circumstances outside the agent; it is determined by the agent's internal states. Other, broader, concepts of content are not the ones on which psychological laws rely.

Now Fodor has become convinced that this won't do: according to him all we have to work with is informational content, which he thinks has to be broad content. So he accepts (2), which, along with (1) and (3), create the problem for this book:

(2) Content is informational, and hence broad rather than narrow, i.e., it does depend on circumstances outside the agent, and is not wholly determined by the agent's internal states.

The problem for Fodor is reconciling (1), (2) and (3); this is what the book is about.

The problems raised by these theses are brought out by the cognitive versions of Frege cases and twin cases.

In a Frege case, there are two ways of thinking of the same thing. They have different causes and different effects in ways that lead to exceptions to putative psychological laws. Suppose, for example, Elwood knows that Cicero was a Roman orator, likes to please his teacher, and has been asked, "Was Tully a Roman orator?" A plausible principle of intentional psychology is that if someone knows that x is P, is asked whether x is P, and wants to please the person doing the asking, he will say that x is P. But Elwood won't do that.

In a twin case, there is one way of thinking of two things. People seem to fall under the same computational laws, without sharing the requisite contents. So Elwood thinks that water is good to drink, and Telwood thinks that twater is good to drink, and these states both lead them to pick up the glass in front of them, bring it to their lips and sip at it. The same computational story linking percept, belief and action applies, but not the same intentional law, for there is no broad content they both believe.

Fodor's solution is that these cases just don't occur very often. That is, in actual circumstances, or anything close enough to them to matter, sameness of computational state will mean sameness of intentional state, and difference of computational state will mean difference of intentional state. That is, there aren't often, and aren't likely to be two ways of thinking of the same thing that are unlinked as the Cicero and Tully ways were in the example, and there aren't often, and aren't likely to be, two things that we think about in the same way as Elwood and Telwood did in that example. Frege cases, Putnam cases, Burge cases and their ilk don't happen very often, and the intentional laws of psychology needn't be perfect.

I don't think this solution is correct. There is not a particularly close tie between informational content and broad content; informational content can be as narrow as one needs. Fodor is overimpressed by the referentialist tradition. It shows nothing more than that folk linguistics and folk psychology have a keen interest in relatively broad content. Absolutely nothing follows from this about the narrowness of content that cognitive science might need, nor the narrowness of content that an informational approach can provide. Whether or not there is anything that should be called absolutely narrow content, contents as narrow as are needed for any particular purpose can be provided by the framework of information. Frege cases and twin cases are not rare events; they occur everyday, pose no significant problem for those who need to describe or explain them in ordinary language. Their possibility is built into the structure of informational content.

2 Everyday Frege Cases

I assume, in the context of Fodor's scholarship, that by a Frege case we don't mean a case Frege actually discussed or thought about, but something that meets some abstract conception of the puzzles he introduced that interests Fodor. I think that means that a Frege case involves a single individual who has two ways of thinking about the same object or property without realizing it. This is not rare.

Suppose I am walking across campus and see someone approaching, someone

whom I eventually recognize to be Fodor. Before the episode begins I have one way of thinking about Fodor, as "Jerry Fodor." I may have been thinking about Fodor's books, his ideas, his boat, etc. As I walk along, I see someone in the distance. At first I think of the person I see as "that man". Then it occurs to me that there is something vaguely familiar. Is it Paul Newman? Is it Bill Clinton? Is it David Israel? No, it's Jerry Fodor. During the interval when I am figuring out who I am looking at, I have two unlinked ways of thinking of the same person, ways we can roughly express as "Jerry Fodor" and "that man". If we construe information broadly, then when I first see Jerry I have the information that Jerry Fodor is at Stanford, for I have the information that *that man* is at Stanford, and that man is Jerry Fodor. But if, right at that point, you distracted me from my recognition process and asked, "Are there any famous philosophers on campus," I would have said, "Not unless you count the regular crew," or something like that, not "Jerry Fodor is," which is what I would certainly say a moment later.

Recognition is a familiar process that resolves a problem we cannot even state if we restrict ourselves to content individuated broadly. And yet recognition is a part of folk psychology, not particularly difficult to describe. How can this be, if we are restricted to broad content?

Or suppose that I am visiting friends in Norway and I step on their bathroom scale. It reads "90 kg" I don't know whether I have lost or gained weight. On the other hand, if it read "200 lbs", I would immediately realize I had gained a few more pounds and become depressed. The effect of the signal in kilograms—its cognitive and hence emotional significance to me—is quite different than that which a signal in pounds would have had. Note that this can be true even if I am pretty familiar with the metric system and weigh things in kilograms for all sorts of purposes, as long as the conversion of weights in the neighborhood of 200 lbs is not completely automatic. Yet the property of weighting 90 kg *is* the property of weighing 200 lbs There are two different symbols, embedded in two different systems, connected with two different ideas, of the same objective state. A Frege case. Not rare, not mysterious, and certainly no problem for a decent theory of informational content.

After we take a brief look at information, I will argue that twin cases, properly understood, are not rare either.

3 Informational Content, Broad and Narrow

Fodor doesn't want to provide us with a theory of content or discuss any details about content at all. It seems fair, therefore, to examine his point of view from the standpoint of the simplest, most straightforward and most metaphysically benign theory of informational content of which I am aware. This is the theory of reflexive and incremental information, developed by David Israel and myself. I'll call it "the incremental theory" for short.[2]

[2] See David Israel and John Perry, "What is Information?" in P. Hanson, ed., *Information, Language, and Cognition* (Vancouver: University of British Columbia Press), pp. 1–19; "Information and Architecture," in J. Barwise, J. Gawron, G. Plotkin, and S. Tutiya, eds.,*Situation Theory and Its Applications, vol. 2* (Stanford

The incremental theory doesn't rely on anything except whatever account of laws, regularities, or constraints one uses for the rest of one's philosophy. The information an event (call it a signal) contains is what its occurrence means about what the rest of the world is like, given a set of constraints that relate the occurrence of signals of this kind to other things. As Hume pointed out, events don't mean anything at all, except relative to some laws or constraints. The basic concept of information carried by a signal σ is what needs to be true for σ to occur, given a set of constraints $\mathcal{C}$. If the constraints are actual, the informational content is really information and must be true; otherwise it may be false, but still quite useful for classifying systems attuned to the false constraints. Someone may think that mushrooms nourish, even though only some kinds do. The state of mind of such a person, as they eat mushrooms, can be usefully classified by the informational content of that event given the constraint: they expect to be nourished.

Strictly speaking, from just this bare Humean concept, all we get is *reflexive* information, that is, information about the signal itself. In Hume's famous example, σ is an occurrence of bread-eating. Relative to the constraint that (eating) bread nourishes (the eater), σ carries the information that the eater in σ will be nourished. We can think of this as pretty narrow information.

Often there are other circumstances we also hold fixed, and hence get broader kinds of information. *Given* that David is the eater in σ, σ carries the information that David will be nourished—information about David, not about σ. This is called incremental information, for it is what is *added* to what is given by the signal.

My perceptual state contains the information that Fodor is on campus, given the fact that the man I see is Fodor. If we fix that fact, the information I pick up perceptually is that Fodor is on campus. If we don't fix that fact, we can still describe the content of my perceptual state in terms of information: that someone is on campus, that someone distinguished-looking is on campus, etc. The less facts we fix, the more narrowly we are describing the content. There are all kinds of informational contents available to characterize signals, including speech and cognitive states, and contents can be as narrow as one needs. A rule of thumb is that if we are trying to explain someone's behavior by reference to the contents of their perceptual and cognitive states, we shouldn't fix facts that are not fixed by their cognitive states.

So, if you know that I have recognized the man I see as Fodor, you can say "John saw that Fodor was on campus". You can say that what I saw was just what the Stanford Daily announced, with the headline "Fodor on campus", for example. The two signals, my perception and the headline, have the same content, relative to quite different constraints and facts. We are forced to broad content, to get at commonalities of this sort.

But to deal with the process of recognition, we are forced to narrower content. I knew at the beginning that the man I was seeing was on campus. That is a content involving me and quantifying over men, and not the same content as the Daily headline

University: Center for the Study of Language and Information), pp. 147–60; and various papers in John Perry's, *The Problem of the Essential Indexical and Other Essays* (New York: Oxford University Press, 1993), especially "Fodor and Psychological Explanations," also written with Israel.

had. When we say that I saw a man on campus, and after a bit recognized that it was Fodor, we are describing a transition from knowing that the man I'm seeing is on campus to knowing that Fodor is.

Understanding exactly how we manage to *describe* this sort of event involves knotty problems in the philosophy of language. But understanding the syntax and semantics of content ascriptions is not a precondition for having a good grasp of how minds handle content, any more than understanding the semantics of the language of everyday physics is a precondition for a theory of billiard ball interactions. The sort of change we are getting at is familiar to us from virtually any system we have for managing information. I have something like a file or a dossier on Fodor that includes a lot about him including his name and what he looks like. I open something like a file on the person I see, and accumulate information in it. When there is a match, I merge or link the files in some way, and the information flows between them; I learn that the person I see is a philosopher, and I learn that Fodor is on campus.

The incremental theory describes the informational content of the perception in terms of the *increment* or addition it makes to some body of information. (Note that the body of information thus incremented does *not* become part of the content. To return to our first example, the broad content of "The author of *The Elm and The Expert* likes to sail" takes as fixed the fact that Fodor authored the book. The increment it adds to what is fixed is *that Fodor likes to sail*. But the worlds in which Fodor likes to sail include those in which he liked it so much that he never got around to writing the book.)

Typically, the body of information we have in mind is that held by the perceiver. In this case, we fix only the facts that are fixed by the agent's cognitive states.

Consider a file drawer with two files about Jerry Fodor, one labeled "Jerry" and the other labeled "Fodor". These sort of duplicate files are not a rare occurrence in systems for managing information about people. (Sometimes there is a good reason to keep more than one file of the same person. A university will typically have a number of files on a given student. The file in the Bursar's Office may contain financial information not open to the student's advisors; the file in the Advising Office may contain details about academic problems of no interest to the bursar.) Then imagine a second filing drawer, just like the first, except that a rubber band has been placed around the two files to indicate that they are of the same person, as one might do after discovering that the Jerry in question was the Fodor in question.

The second file drawer contains information that the first one does not. We can't get at this information if we stick to broad content.

Consider, for example, the possibilities that meet the conditions imposed by the first drawer, assuming a referentialist account of "Jerry" and "Fodor'. They are exactly the possibilities that meet the conditions imposed by the second drawer because, sad to say, there is only one Jerry Fodor. If we fix the facts about who the files are of, there is no way to get at the possibility that they are of two different people, the possibility that the rubber band in the second drawer rules out, and so no way to get at the difference. Referentialism imposes upon us an inappropriate degree of breadth for describing the informational content of the file drawers. Slavish referentialism wouldn't allow us to explain why, say, the money owed Jerry isn't sent to him, given that his address is right

there in the Fodor file.

But there is nothing about *information* that limits us to such broad content. Informational content allows a graceful retreat to any number of narrower contents, that bring out the differences between the drawers. Suppose we fix the fact that the "Fodor" file is of Fodor, but do not fix the fact that the "Jerry" file is. With those things fixed, the conditions the first drawer imposes can be met even if the files are of different people, but the conditions imposed by the second drawer cannot be.

Fodor also sees informational content as limiting in another way that is puzzling. He seems to equate informational content with the information a signal contains about its *causes*, which leads him to see a conflict between informational and functional approaches. But the information contained by a signal relative to a constraint can be about the future or the past, about effect as well as cause, about behavior as well as perception. There is no particular reason that those who find their content in information cannot also be functionalists.

Suppose you believe that Fodor doesn't like to sail. This is misinformation about Fodor, since he does like to sail. Still, your being in that state carries information about you and how you will act towards Fodor, given that you have the means to recognize him. If you like him, you won't ask him to go sailing with you. A state of mind, like a state of anything, carries all sorts of information relative to different constraints and different fixed facts. Why do we characterize such states in terms of broad content, content involving Fodor, say? Because just as our cognitive capacities evolved to enable us to carry information about individuals with whom we interact, folk psychology evolved to enable us to efficiently characterize individuals who are putting those cognitive capacities to use. Describing mental states at work in terms of broad content rather than narrow is as natural as describing a wrestling match in terms of what the participants are doing to each other rather than anatomically.

4 Everyday Twin Cases

When we talk about broad and narrow content, we have in mind a type of signal, some constraints relative to which these signals do or would provide information both about causes and effects, and nested sets of facts, thought of as less and less restricted, that provide broader and broader content. Let f and f' be such sets of facts, where f is the more and f' the less restricted, that is $f \subset f'$. In a Frege case two signals have the same content relative to f' while having different contents relative to the more restricted set f. In a twin case, two signals have the same contents relative to some more restricted set of facts f, but different contents relative to some less restricted set f'.

In the literature, twin cases usually incorporate some kind of unlikely duplication of detail, as with Elwood and Telwood, and so of course are unlikely to occur. There is nothing essential about this massive duplication; its role is purely dialectical. The twin case is supposed to convince us of the possibility of differences in broad content in spite of sameness of internal state. To be convincing, one has to make the internal states similar in any ways one's interlocutor might deem relevant. If one's interlocutor

is a holist, worships holists, was taught holism in graduate school, and the like, one needs massive duplication, so that one's interlocutor cannot trace difference of content to internal difference.

But the fact that massive duplications are rare doesn't help Fodor's case. He's not a holist, and has become a fan of broad content determined in part by external circumstances. There is nothing in his view that suggests that twin cases will be rare if by twin cases we mean differences in broad content across individuals with *relevantly* similar cognitive states.

5 Conclusion

At the heart of the concept of information are general principles about the way things work, in virtue of which specific signals at different times and places and connected to quite different things can carry information. The possibility of Frege cases and twin cases is built into the very structure of information. The same information can be carried by signals of quite different kinds and quite different causal roles relative to different facts and/or constraints: Frege cases. The same signals, connected to different objects, can carry different information: twin cases. If we include, among the differences that can make a difference in the incremental content of similar signals, the other signals they combine with, then systematic semantics is simply the study of twins.