

The Architectural Nonchalance of Commonsense Psychology

RADU J. BOGDAN

1. The Issue

Eliminativism assumes that commonsense psychology describes and explains the mind in terms of the internal design and operation of the mind. If this assumption is invalidated, so is eliminativism. The same conditional is true of intentional realism. Elsewhere (Bogdan 1991) I have argued against this 'folk-theory-theory' assumption by showing that commonsense psychology is not an empirical prototheory of the mind but a biosocially motivated practice of coding, utilizing, and sharing information from and about conspecifics. Here, without presupposing a specific analysis of commonsense psychology, I want to challenge a key implication of the 'folk-theory-theory' assumption to the effect that commonsense psychology is committed to a definite architecture of the mind.

It is this architectural commitment that invites the elimination of commonsense psychology by connectionism (Stich 1983, ch. 11; Stich 1988; Ramsey, Stich, Garon 1990) or its vindication by the language of thought or syntax-over symbols view, which I will call syntactic computationalism (Davies 1992; see also Fodor 1975, 1981, 1987; Pylyshyn 1984). Since the study of mental architecture is scientific, why bother with commonsense psychology?

Why are its alleged architectural commitments important as a basis for elimination or vindication? Three reasons spring to mind. One is that both eliminativism and realism use commonsense psychology as a stick to prop their scientific candidate and beat on its opponent. For eliminativism, in particular, if connectionism is right about mental architecture, not only is commonsense psychology wrong and eliminable, but so may be its scientific sponsor, syntactic computationalism. This is because commonsense psychology, which has been around for much longer than syntactic computationalism, is often viewed by friend and foe alike as a source of inspiration for syntactic computationalism.¹

Another reason is that commonsense psychology appears to provide an apriori conception of our cognitive selves, as thinkers, which emphasizes a formal rationality that can be implemented only by the architectures envisaged by syntactic computationalism. Eliminating commonsense psychology on architectural grounds is thought to explode this apriori conception that eliminativists find wrong and counterproductive. Finally, there is the ability of commonsense psychology to explain and predict cognitions and actions successfully. The thought naturally occurs that the success of this ability must

¹ For intentional realists like Fodor and Pylyshyn the commonsense framework is the inspiration for the cognitive-scientific framework: see Fodor 1975, ch. 1; 1981, Introduction, chs. 4, 7; 1987, Preface, ch. 1; Pylyshyn 1984, ch. 1. Ramsey, Stich, Garon (1990) explicitly talk about "one sort of psychological model that exploits and builds upon the posits of folk psychology" (p. 500), and later hold that "in the psychological literature there is no shortage of models for human belief or memory which follow the lead of commonsense psychology in supposing that propositional modularity is true" (p. 506). There are many other examples of the cognitive science-common sense connection in the philosophical and foundational literature.

have an architectural grounding since its explanations and predictions are causal.

I find none of these reasons for the architectural significance of commonsense psychology compelling. The last two will be taken up later on, as we speculate that commonsense psychology is likely to view a thinker functionally, as a problem solver, rather than architecturally, as a symbol manipulator; and that its explanatory success may rest on its architectural nonchalance rather than commitment. The first reason, having to do with using commonsense psychology in the battle between connectionism and syntactic computationalism, is a bogus that had better be disposed of right away.

In spite of the well documented influence of commonsense psychology on the philosophy of syntactic computationalism, the major inspirations of the latter come from other directions. I am thinking of Chomsky's theory of grammar, early psychological models of concept formation, memory, and thinking, computer simulations of cognition, computer vision, and the reality of computers. As the professional literature testifies, it is in terms of such phenomena that the merits of connectionism and syntactic computationalism are judged. Significantly, after handwaiving in the direction of commonsense psychology, the archintentional realist, Jerry Fodor, gives his most forceful arguments for syntactic cognition in terms of compositionality, computation, systematicity, and productivity, and illustrates them with structural descriptions of visual scenes and parsing trees (Fodor 1975, chs. 3 and 4; 1987, Appendix). These are hardly the sorts of features that commonsense psychology cares about, and are precisely the sorts of features that syntactic computationalists

think connectionism cannot explain (Fodor and Pylyshyn 1988). Similarly, when Ramsey, Stich, Garon, or Davies, discuss connectionism pro and con, the examples are memory retrieval, formation of tenses, and phonetic learning -- again, hardly the stuff about which commonsense gets excited or excessively opinionated.

This matter being clarified, the question we turn to is whether commonsense psychology has architectural commitments. I think it hasn't, which is why I want to show that the architectural elimination (or vindication) of commonsense psychology is a misguided project. My plan is this. I examine in the next section the notion of propositional modularity proposed by Stich (1983), and developed by Ramsey, Stich, and Garon (1990), suggest that it allows different architectural readings, and then force it into a dilemma by arguing that, on an implementational reading, propositional modularity may be relevant to commonsense psychology but has no architectural bite, while on a stronger reading, it has the opposite effects, bite but no relevance. Then, in the following three sections, I examine the evidence (in the form of apriori notions, introspection, and attributions of propositional attitudes) invoked for the architectural commitments of commonsense psychology, and find it inadequate. I conclude with a speculation about how the architectural nonchalance of commonsense psychology bears on its explanations. To simplify the terminology, I describe as archphilosophers all those who think that commonsense psychology has architectural commitments, whether they are eliminativists, realists, or of other stripes. My argument, here, is not that they are wrong (which they are) but rather that they have not made a convincing

case for their position.

2. Degrees of Architectural Involvement

Propositional modularity is the notion that mental states are (i) functionally discrete, (ii) semantically interpretable, and (iii) play a causal role in cognition and behavior in virtue of (i) and (ii). Stich (1983, pp. 237-242) and Ramsey, Stich and Garon (1990, pp. 504, 506) characterize this notion in the commonsense terms of beliefs, memories, and other propositional attitudes, and take it to have architectural import. For our discussion, it suffices to say that the notion of cognitive architecture captures the resources that are systematically responsible for the form and operation of representations. The resources include basic components of representations (e.g., symbols or neural nets) and basic programs that handle representations. For syntactic computationalism, these are resources that store, retrieve, and manipulate symbol structures according to their form. This indicates a certain "style of processing" (Fodor and Pylyshyn, 1988, pp. 5, 10) in which programs articulate simple symbols into complex structures, and combinatorially manipulate such structures according to their composition. Connectionism contrasts this "style of processing" with its own, based on a nonsyntactic and nonsymbolic architecture of activation networks.

The notion of propositional modularity allows several degrees of architectural involvement. The most minimal is that of an architecture that merely implements propositional modularity. This is the implementational sense

in which elementary particles implement thought or paper transactions implement money exchanges. Since in its talk of beliefs and desires, commonsense psychology is said to posit states that are functionally discrete, semantically valued, and causally effective, it must be committed to some unspecified architecture that implements propositional modularity. What the implementational commitment does not entail is that the architecture itself is propositionally modular. To get that entailment, we need an argument for a stronger architectural involvement.

At this point, we should note in anticipation, a dilemma is beginning to take shape. On the one horn, the minimally implementational degree of architectural involvement may be the closest the notion of propositional modularity comes in approximating the architectural commitments of commonsense psychology. Unfortunately, this is too shallow and trivial a commitment to be invalidated by connectionism (or vindicated by syntactic computationalism). On the other horn, any stronger commitment might allow such an invalidation (or vindication) but at the price of misrepresenting commonsense psychology.

Eliminativism and realism need a stronger architectural involvement on the part of commonsense psychology if their favorite cognitive-scientific paradigm is to yield elimination or vindication. According to Ramsey, Stich and Garon (1990), commonsense psychology appears to have such a stronger involvement insofar as it entails that beliefs and desires are acquired or lost one at a time, or that the same belief is involved in various inferences, or that out of several possible beliefs only one was causally active in generating an action, or the like. These entailments point to properties that only a propositionally

modular architecture could have, and to processes that only such an architecture could run. In addition, commonsense psychology needs a stronger architectural involvement to satisfy the conditions on propositional modularity. Specifically, the third condition, on causation, could not be met unless the functional units were causally efficacious in virtue of their semantic values. That causal efficacy requires appropriate architectural resources.

How strong should this new involvement be? The line taken by Stich and his collaborators points to a rather weak involvement that bears on mental sentences as units, and not on their internal form. This cannot be good enough for the archphilosopher. Left unanalyzed, mental sentences have no psychological value, not only because their formation and transformation are left in the dark, but also because, as a result, we have no understanding of the causal mechanics by which mental sentences play the functional roles required by propositional modularity. Mental sentences are neither semantically nor causally potent merely in virtue of having truth values. Semantically, they represent because their parts do; causally, mental sentences could not represent, and could not generate other sentences, unless they are manufactured and recombined out of parts. Psychologically, it is precisely the internal composition of a mental sentence that makes its semantic value causally efficacious.² This is an axiom of syntactic computationalism. I am not

² Ramsey, Stich, Garon (1990, p. 506) write that the propositionally modular psychological models inspired by commonsense psychology view "beliefs or memories as an interconnected collection of functionally discrete, semantically interpretable states which interact in systematic ways." I submit that no plausible account of the interconnectedness and systematic interaction of beliefs and memories can avoid appeal to the internal composition of sentences. Stich himself admits that the "unstructured format" of sentential

saying that this is the right axiom, or that commonsense psychology is committed to a syntactic architecture. I am saying only that a sentential reading of propositional modularity leaves the route from semantics to causation gappy if not mysterious, thus preventing commonsense psychology from coherently satisfying all the conditions on propositional modularity.

Two further facts militate against a weak sentential involvement. One is that the ordinary attributions of propositional attitudes are notoriously sensitive, both semantically and causally, to the internal constituents of the attributed contents. We know that beliefs fail to retain their semantic value and causal efficacy if, when equivalent parts of their contents are substituted, the equivalence is not recognized by the believer. The archphilosopher would want the architectural commitment of commonsense psychology to reflect this intrasentential sensitivity. Another fact is that connectionism is interesting generally, and as an elimination ploy, in particular, precisely to the extent to which it is an alternative to a syntax-over-symbols architecture, and much less interesting if merely suggesting nonsyntactic mechanisms for hopping from one sentence unit to another.

An intrasentential reading is therefore needed for the alleged architectural commitments of commonsense psychology to meet coherently all the

models has a hard time locating and tracking the right information routes in memory recall and other tasks (Stich 1983, p. 239), after having agreed with language of thought theorists that propositional (i.e., semantic and nonmental) theories of belief misrepresent commonsense psychology (op.cit., p. 29) I am not sure that is such a misrepresentation, but the fact is that a mental sentence theory comes close to have the same psychological irrelevance as a nonmental propositional theory of attitudes.

requirements of propositional modularity, and thus be subject to elimination (or vindication). This is the line taken by Davies (1992) with respect to connectionist eliminativism, and I think it is the right way to go. Yet I hasten to add that the conclusions of this paper also apply to the weaker reading of propositional modularity. I call this stronger reading syntactic modularity. The question, then, is whether commonsense psychology has a strong architectural commitment to syntactic modularity, as archphilosophers (should) claim. If so, what sort of commonsense psychology do they have in mind, and what is their evidence for the claim? These are the questions we turn to next.

3. The Evidence

As understood here, commonsense psychology is not what ordinary folk think about the mind, whatever that may be. As I see it, commonsense psychology is a universal competence, possibly innate and surely matured at an early age, which is constituted by a repertory of concepts and attribution schemas that are employed naturally and unreflectively to interpret and predict cognitive, emotional, and behavioral conditions (Bogdan 1991). People exercise this competence spontaneously and successfully, irrespective of what they happen to think about the mind. People often have the silliest views about the mind (as the history of philosophy and folklore abundantly testifies) and yet succeed in interpreting their conspecifics by using commonsense psychology.

In a word, commonsense psychology is one thing, false consciousness about the mind or about commonsense psychology itself another thing. Neither

form of false consciousness is good evidence for the nature and operation of commonsense psychology, and neither plays any role in our story. The argument here is that archphilosophy has not shown that the commonsense concepts and attributions entail or evidence a commitment to syntactic modularity. The argument is not whether commonsense psychology is right or wrong in its architectural commitments, but whether archphilosophy shows that it has them. (I deny that it shows that.) The argument is concerned only with the evidence that motivates the archphilosopher to infer to the alleged architectural commitments of commonsense psychology. (I deny that the evidence motivates such an inference.) Again, the argument is not about what ordinary folk think of, or infer from, that evidence but rather about how the archphilosophers reconstruct the evidence to substantiate their charge.

The archphilosophical position holds that the functionally discrete, syntactically organized, and semantically interpretable mental states with causal roles in cognition are the propositional attitudes of commonsense psychology. It also holds that the commonsense conception reflects architectural commitments to syntactic modularity. How do archphilosophers know these truths? Not scientifically. We do not know scientifically how commonsense psychology works. Aside from very recent studies in animal and developmental psychology, no systematic inquiry has been mounted to find out the rationale, nature, and operation of commonsense psychology. But there is no need for such an inquiry, the archphilosopher would say, smiling knowingly. For we have not only apriori and introspective evidence, but also the evidence of the attributions of propositional attitudes, to conclude that commonsense

psychology is committed architecturally to syntactic modularity. These three sources of archphilosophical evidence are worth scrutiny.

THE APRIORI EVIDENCE. Commonsense psychology must be committed to syntactic modularity in an apriori sense because this is what the commonsense concept of thinker means (Davies 1992). To be a thinker is to instantiate a syntactic architecture. My reply is that even if we grant that the commonsense concept of thinker could be determined apriori, it does not follow that the concept entails syntactic modularity.

For a softening up, consider this possibility. We could conceive of a thinking architecture with these properties: (i) it is not syntactic; (ii) it executes most of its cognitive tasks in nonsyntactic formats; yet (iii) for a definite class of tasks (e.g., logical, linguistic) it emulates syntactic modularity by "creating artifacts in the form of physical representations that we can manipulate to get answers to very difficult and abstract problems" (Rumelhart, Smolensky et al, 1986, p. 44). This (connectionist) position does not have to be right, only conceivable. Since it is, thinkerhood does not entail a syntactic architecture. Still, you would retort, this is not commonsense conceivability. In letter it is not, but in spirit it may be. As we will note shortly, in discussing introspection, we are internally aware of free and unorganized perceptions and ideas that get trimmed, firmed up, and syntactically regimented once they are formulated linguistically. This realization allows us intuitively to conceive of a narrow and disciplined architecture, on the top of a much larger and undisciplined one, getting to work intermittently when certain types of cognitive tasks must be carried out.

A more direct objection to the apriori argument is the following. It is not clear what we mean by an architecturally syntactic thinker. Is a syntactic architecture a necessary and sufficient condition, or only necessary? It is not logical niceties but levels of explanation that matter here. From an apriori stance, a thinker may be a means-ends reasoner or problem solver rather than a syntactic modulator. If a person or animal appears to deliberate and choose, commonsensically we are inclined to count her a thinker, without knowing or bothering about how she does it. We take a similar stance toward intelligence as a measure of how well or quickly problems are solved rather than as a description of the cognitive means employed. Pressed with examples and theory, the commonsense psychologist might admit that syntactic modularity is a way of implementing means-ends thinking. Syntactic modularity would then be a necessary but not sufficient condition for being a thinker. In that case, from the standpoint of commonsense psychology, logic and language appear as syntactic implementations of thinking. We think (i.e., solve problems, pursue goals) by means of logic and language programs, for definite tasks. As with other implementations (e.g., atomic forces ensuring the solidity of objects), commonsense is wisely agnostic about how the mental logic and language work architecturally. The implementational sense of architectural involvement once again advertises its plausibility.

INTROSPECTION. What about introspection, the inner phenomenal prop of commonsense psychology? When treated commonsensically, doesn't it show that thinking is syntactically modular? Let us see. What exactly are we

introspecting from the stance of commonsense psychology? What does the latter enable us to grasp and conceptualize from the passing show of our inner mental life? Certainly neither programs nor underlying mechanisms. These cannot be introspected; and even if they could, we wouldn't get them right anyway. As psychologists have known for decades, introspection is biased, unrealistic, and unreliable.

As far as introspecting thinking is concerned, I see two alternatives. One is a passive, kibbitzing introspection, a diffuse awareness of what we are doing mentally. This mode of introspection often fails to deliver the "structured thought" required by syntactic modularity. Spontaneous thinking is a typical example. I do not know about you, but I introspect and commonsensically identify little structure when thinking spontaneously. I introspect (mostly) free, anarchic, dynamic, and diffuse associations of barely formulated and unfocused ideas, assaulted and intermittently displaced by other sensory and memory inputs. If, Heaven forbids, my naive introspection of my spontaneous thinking were generalized as commonsense psychological evidence of functional architecture, connectionism would win hands down.

The other alternative is our introspection of linguistically regimented outputs and of socially disciplined forms of argument and attribution. We so introspect when we talk to ourselves, or think aloud, or express our thoughts publicly, or engage in attributions, including self-attributions, of propositional attitudes. Am I not aware and conceptualizing commonsensically (right now, as I develop this very thought, talking sotto voce to myself) that I think in forms that are syntactically modular? I might be (OK! I am) but only of the evidence of

the outputs of the operation of my mental architecture. Informed by commonsense concepts, my introspection does not reach the architecture itself or its modus operandi, nor do those concepts give me any reason to infer from the introspected data (disciplined outputs) to the syntactic modularity of the architecture. As we noted already that, the output evidence is compatible with a different (connectionist) architecture or else uninformative about architecture.

ATTRIBUTIONS. Yet surely, the archphilosopher would retort, the strongest case is made by propositional attitudes and their attributions. When I attribute to you the thought that cats are more fun and less noisy than barking dogs, isn't my commonsense attribution assuming the presence in your head of data structures made of grammatical and logical parts, functionally discrete, semantic, and, when attitudinized, causally efficacious in virtue of these features? Isn't this a commitment to an architecture responsible for such features? When we take one to have acquired a belief or lost a desire, isn't our commonsense conceptualization assuming that semantic and functionally discrete units are causally active or inactive, respectively, in virtue of their architectural instantiation or lack of it?

The archphilosophers can make their case only if they have a solution to (what I call) the gap problem. There is a gap to be bridged between the attributions of propositional attitudes, on the one hand, and the data structures and processes, and their architectures in the head, on the other. The solution to the gap problem must consist in a principled and systematic mapping of

propositional attitudes into syntactic architectures. The gap problem is at the heart of the debate about the elimination (or vindication) of commonsense psychology, so it deserves some discussion. In the next section we examine the general dimensions of the gap problem, and then consider why it may have no solution.

4. The Gap Revealed

One solution to the gap problem would be simply to assume that there is a mapping from propositional attitudes to architectures. Some archphilosophers may have bought this solution unconsciously. Since we are interested in arguments, not fictions, we move on. One serious archphilosophical argument might go as follows. A reason why people can be said to have propositional attitudes in their heads is that the commonsense psychological discourse so portrays them. The attribution language portrays people as possessing beliefs and desires, and also portrays beliefs and desires as doing causal work. This is not enough. The mental form of the attitudes must be also established, and that form must be syntactically modular. But we know what the mental form is (the argument continues) because the attribution discourse displays it.

The attribution discourse does display that-clauses which describe contents, and attitudes which describe the roles of the contents. The question is to place the contents and attitudes in the mind and show three things about their mental realization: (1) that they describe data structures and processes instantiated in the mind; (2) that the instantiations obey constraints that

reflect the conceptual, logical, and grammatical features of the attribution sentences; and (3) that the constraints have causal efficacy in virtue of what they reflect. If the archphilosopher can show all this, then the architectural commitment is just around the corner, since the mental instantiations of propositional attitudes would not be possible without architectures that embody causally the required conceptual, logical, and grammatical constraints. Can the archphilosopher get the instantiations and the constraints in a principled and systematic manner? If he can, I am off to Carcasonne and a life of medieval pleasure.

To see what the archphilosopher's task is, let us first dispose of a few dummies. Some are familiar, others less so. Suppose I (the generic commonsense psychologist) attribute to David (the generic subject) the belief that the sand is hot. For the attribution to work, none of the following need be true: my natural language need not match David's; David need not speak any language, nor need he think about the hotness of sand, and what follows from this, the way I do; David need not explicitly formulate his belief, publicly or to himself; my concepts, experiences, inferences, or intellectual powers need not be similar to David's; David could be a small child or a chimp, and still worthy of the same attribution; I may make the attribution even when the evidence for it changes; the evidence may draw on noticing David's face or his behavior or on what he says, in so many words, or in other words, and so on. The archphilosopher is advised to begin to worry: How can the attribution stay put with so much variation around?

Perhaps less obviously, there are other facts that should not be appealed

to in a solution of the gap problem, even when we consider only interpersonal attributions involving adult language users. The solution should not rely either on the fact that a natural language may have a syntactic architecture; or that a good deal of our cognition, including our cognition of others, is permeated by language; or that we express publicly our cognitions in a natural language, and use it to formulate our attributions of propositional attitudes. The syntactic modularity of the language faculty is neither available nor important to commonsense psychology. The latter has access solely to linguistic performances, not to their underlying competence. (This is why the commonsense psychologist may have a successful attribution practice while having a false consciousness about the practice and its success conditions.)

A natural language is a tool used in attributions. Whatever the properties of the tool, they do not show a commitment to the syntactic modularity of our mental architecture. An independent argument to that effect is needed. Since it is the surface (performance) structures of our language-based attributions of desires and beliefs that must be analyzed to figure out the operation and commitments of commonsense psychology, we cannot simply project those structures into the mental architecture and expect systematic mapping. Not, if we do not want to beg the whole question. When asking David whether he believes that the sand is hot, and he answers that he does, or volunteers in English the belief that the sand is hot, or puts it this way to himself, the archphilosopher is still not entitled to conclude (without a further argument) that the belief thusly attributed or self attributed reflects a commitment of commonsense psychology to how David's cognition works architecturally.

Language and cognition are one thing, commonsense psychology another. We cannot map systematically the attributions of the latter into the architectural innards of the former merely because language and cognition (syntactically modular as they may be) are presupposed in the exercise of commonsense psychology. The atomic architecture of the world is presupposed in our commonsense physics, yet we would not say that the latter maps systematically into the atomic architecture of the world. It is not simply that the commonsense concepts are more superficial than the scientific ones (which they are); their jobs are different. Most commonsense physical concepts are partly functional, whereas the concepts that physical science deploys are not. Commonsense physical concepts enable us to handle objects, navigate around them, form behavioral expectations about them, and the like. Their degree of intellectual penetration into the nature of physical things is commensurate with their practical functions.

And that is precisely the point! Commonsense concepts are functional since practical. So the question to ask is what is the function of the commonsense psychological concepts. It surely would beg the question to say that their function is to map systematically into the architecture of the mind in order to give an accurate description of someone's mind. Even if that is so (which I so much doubt), the key problems are still with us. Why would common sense do that? How successful is the mapping and the resulting description? And do they reveal an architectural commitment to syntactic modularity?

Let me come at it from another direction. Suppose the mind is syntactically modular, and commonsense psychology is committed to this truth. It still does

not follow that such a commitment owes anything to attributions of propositional attitudes. For the supposition is compatible with the attributions having a job other than that of describing the architecture of the mind or a job that does not require such a description and does not presuppose a syntactic architecture. For example, the attributions may have the job of conveniently summarizing and editing what the subject is saying or how she is behaving; or that of simulating the subject's cognitive condition from the perspective of the attributor; or that of providing the attributor with information about the subject and her world; or something else. These are lively possibilities pursued in the literature. If any or several of these possible jobs are the right ones, it is no longer clear what it is in the head of the subject that the attributions describe and presuppose.

Let us now twist the supposition that the mind is syntactically modular even further. As far as cognitive science is concerned, it is possible that mental architectures generate representations that are functionally efficacious in a discrete way without being semantically interpretable. According to methodological solipsism, none of the structures and processes posited by cognitive science should be construed in semantic terms, if the causal efficacy of those structures and processes is to remain intelligible. But then, the notion of syntactic modularity needed by eliminativism is amputated since the semantic component, to be left out, is one which commonsense psychology cannot do without. If anything, commonsense psychology would sacrifice the other components of the syntactic modularity foisted on it, which happen to be the real architectural ones.

Methodological solipsism may of course be either wrong, or else applicable only to occurrent representations, not to the durable architectures that generate them. In that case, the mental architecture can still be given a semantic reading from a cognitive-scientific standpoint. That, however, would not be the standpoint of commonsense psychology. A cognitive-scientific account of the semantics of cognition could invoke evolution to explain sensory-motor architectures, or ecological and optical laws to elucidate the architecture of vision. That is not the game of commonsense psychology. The semantics the latter is after is ephemeral and tracks representations in a context. The interests do not match: the ephemeral and context-sensitive semantic representations that cognitive science is either solipsist about, or indifferent to, are those that interest commonsense psychology, whereas the durable architectures that produce those representations, and the generic semantics the architectures acquire by evolution, are of interest to cognitive science and beyond the reach of common sense.

Such a divergence of interests may bear on the other components of syntactic modularity. It may turn out that what the archphilosophers classify as the causal roles of propositional attitudes has little in common with the mental causation elucidated by cognitive science. It may be that the functional roles posited by common sense are not meant to track actual causal paths in the mind. Suppose that when I predict that David will not step on the beach, given what I take to be his belief that the sand is hot, I am extrapolating from my experience. The latter does not presuppose a mental architecture or, in particular, specific causal paths from representations to behavior. My experience

provides a frame of reference in terms of which commonsense psychology allows me to plot, perhaps simulate, David's linguistic and behavioral performances. The plotting connects some nodes (beliefs, desires) with others (actions), and thus extracts the information of interest. Although the plotting would not work unless there were some causal relations between David's representations and his actions, nothing in our supposition indicates that those unspecified causal relations need match systematically well-defined architectural paths of a syntactic or some other form. To return to our leitmotif, the commonsense commitment to an unspecified mental causation is on a par with its implementational degree of architectural involvement.

5. The Gap Illustrated

So far I have voiced a general skepticism about the possibility of a systematic mapping of commonsense attributions of propositional attitudes into mental architecture. I want now to illustrate my skepticism in more intuitive and specific terms.

The attitude attributions are notoriously flexible and versatile. I can say that she saw the boat, or saw that it was a boat, or saw the boat approaching, or saw that the boat was approaching, and so forth. Suppose I attribute these perceptual attitudes on the basis of her behavior, my perception of her, and the scene in front of us. My attributions keep redescribing the content of her perception, first, in terms of the target of her vision (object = boat; fact = that it is approaching), and then aligning her attitudes to the type of target (object

perception = seeing the boat; propositional perception = seeing that the boat is approaching).

These banalities illustrate what our perceptual attributions explicitly do. What about their implicit commitments to the visual representations she forms and the architectures involved in their formation? How should we construe these commitments? As a generic commonsense psychologist, I am committed (in the case under discussion) to describe her perception externally, or *de re*, in terms of items of interest in her and my visual range (boat, approaching movement, and so on). I also have the causal expectation that she gets the information by vision, that her perceptions animate her thoughts and may lead to action. What I know of her vision is (more or less) what I know of mine: that in keeping my eyes open I see scenes populated by all sorts of objects and events, which I name in ordinary language and classify with the concepts of common sense.

Archphilosophers need a stronger reading for their mapping argument than that presupposed by my commitments. Two steps are needed for that reading. In our example, one step must go from my attribution to her visual images and perceptual attitudes; the other step must go from the latter to her visual architecture. I do not think these two steps can be made. We cannot expect the visual images, with their wealth of analog detail, to be matched systematically by their abstractive linguistic descriptions. What the attributions describe (boat, approaching motion) are not the categories of the visual fields (shapes, volumes, motions, colors) picked up by visual images.

We saw that some attributions of perceptual attitudes are propositional,

others object-directed. These types of attributions can neither individuate nor be committed to distinct forms of vision, let alone distinct visual architectures. There is no evidence that we have a vision for objects and one for facts, with distinct sets of images, and no evidence that common sense would presuppose anything like this. Our commonsense attributions posit a distinction without an architectural difference. It is scenes that we visualize, although our recognitional, conceptual, and judgmental programs decompose them into objects, properties, facts, and events. But our attributions ignore the architectural division of labor involved in perception and in its recognitional and judgmental uptake, and focus instead on the gross semantic features (typically, facts) of the scene.

This diagnosis extends to other forms of cognition. Our doxastic or conative attributions are made in both the propositional format (e.g., she believes that he is right) and the object-directed format (e.g., she believes him). Again, we cannot infer from these attributions that there are distinct representations and architectures for object-directed and for propositional belief. Nor do we have grounds to suspect that, in using such attributions, commonsense psychology is committed to the psychological reality of distinct architectures. The proposition/object and the de dicto/de re distinctions are distinctions without architectural differences (Bogdan 1986; 1991). The attribution-architecture mapping sought by the archphilosophers is still elusive, and the gap as wide as ever.

6. The Wisdom of Nonchalance

Let me conclude with an indirect argument for the architectural nonchalance of commonsense psychology. We noted that commonsense psychology is thought to have architectural significance because of its successful explanations. It is widely assumed that these explanations are causal, and so it is inferred that there must be an architectural basis for their success. I think this inference has got things upside down. To the extent to which commonsense psychology is successful at explanation, it is more likely to be because of indifference rather than commitment to specific architectures. If commonsense psychology were strongly committed to a specific architecture, it would not be as successful as it appears to be. For, at least on one influential analysis, a strong architectural commitment would be an obstacle to, rather than a basis for, successful explanation. This, then, is my indirect argument: since it is successful at explanation, and a strong architectural commitment would be inimical to a successful explanation, commonsense psychology is unlikely to be strongly committed architecturally.

The idea that a strong architectural commitment would be an obstacle to successful explanation draws on Jerry Fodor's (1983) infamous thesis that we know how to think cognitive-scientifically about dedicated (vision, grammar) architectures, but not about undedicated central (thinking) architectures.³

³ There is intriguing research in recent evolutionary psychology showing that even the central programs for reasoning and planning may have been evolutionarily installed first as dedicated and domain-specific architectures and then pirated for other functions (see Cosmides and Tooby 1989). I am going to ignore this development in what follows. If pressed, I would argue that the central presence of dedicated and domain-specific programs does not necessarily exclude the properties of reasoning Fodor is talking about.

Since commonsense psychology is successful precisely in the area of central cognition, where cognitive science is not, it appears that common sense knows something about the central architectures of the mind that cognitive science doesn't. I read the same point differently, that is, nonarchitecturally.

I regard Fodor's thesis as having considerable merit IF a hard core cognitive-scientific theory, that is, one with the strong architectural commitments of syntactic computationalism or connectionism, is adopted as a theory of the central programs (thinking, planning, and so on). The adoption recommends that we conceptualize and explain thinking the way we conceptualize and explain vision or grammar, that is, as production of representations by appropriate architectures. If we do that, then a number of strange properties such as the isotropy and Quininity of central cognition that Fodor is talking about appear intractable to cognitive science. That intractability may be due to a miscalculation that hard core cognitive science makes and commonsense psychology doesn't. The miscalculation is what our indirect argument is all about.

The miscalculation is this. Syntactic or connectionist architectures are constituted by programs that produce and manipulate representations. These are production programs. Hard core cognitive science is committed strongly to production programs. Thinking, by contrast, utilizes representations for a certain job. Thinking is run by utilization programs that presuppose, but are not the same as, production programs. The latter operate as implementers of the former. There are two senses of implementation here: one is the 'raw material'

sense, the other the 'underground work' sense. The former sense is illustrated by vision programs producing the raw material (visual image) to be utilized by perception programs and mapped into belief and thinking programs. The underground work sense is illustrated by the language of thought or any other form of an internal mentales code which translates and runs higher level programs such as speech acts or problem solving programs. The difference here is between a programming language and particular task programs defined and run in its terms. The programming language is a formal implementer. So is the language of thought vis-à-vis the thinking program. The miscalculation made by the hard core theories is not seeing that, viewed in either sense, the production programs cannot sufficiently constrain and explain how the utilization programs work. Isotropy, Quinity, and other global properties of thinking are symptoms that the central architectures of the mind execute utilization, not the implementing production, tasks.

As with any implementation, from its angle and in its terms, we cannot explain what is implemented. Chemical reactions implement but do not alone constrain and therefore explain genetic programs. Closer to the cognitive home, grammar implements but does not fully constrain and therefore explain speech acts. The grammar programs execute production tasks, the speech acts programs execute utilization tasks. In terms of this comparison, commonsense psychology is interested in speech and thinking acts, as utilizations, not in their grammatical and computational implementations. ⁴

⁴ I develop this line of argument further in my (1983; 1988b; and 1994).

Commonsense psychology demonstrably cares about thinking and other forms of utilization. If commonsense psychology were committed strongly to implementing production architectures, as eliminativism and intentional realism assume, then, on my diagnosis of Fodor's thesis, its successful explanations would be a miracle. By contrast, its architectural nonchalance frees commonsense psychology to tap other resources for explanation which hard core cognitive science has either ignored or not yet figured out.

If commonsense psychology presupposes anything about its subjects, it is that they are agents and informants with the minds of problem solvers and deliberators (hence utilizers of representations) rather than of syntactic modulators, symbol pushers, or network activators (hence producers of representations). I am not suggesting that commonsense psychology is a theory of central cognition as problem solving or deliberation. All I am suggesting is that in its practical and unreflective ways, which remain to be explained, commonsense psychology manages to avoid the miscalculation I was trying to identify.

Since commonsense psychology is good at interpreting centrally mediated cognition and behavior, many cognitive scientists are tempted to emulate its wisdom in a strong architectural sense. That has been the policy recommended by intentional realism on behalf of syntactic computationalism. It would not surprise me to hear that somewhere (probably in California) people are having visions of a commonsense psychology that is congenial to the connectionist agenda. (Since it cannot be the eliminativist San Diego, it must be the new age Berkeley.) I hope this paper begins to show why that architectural emulation

would not be advisable. If commonsense psychology provides any scientific inspiration, it should come from the realization that the thinking mind is more than individually shaped, and is run by architectures other than the production ones. Commonsense psychology may be wiser than it seems about what makes humans intelligent, socially informative, and goal-directed. ⁵

Department of Philosophy

Tulane University

New Orleans, LA 70118

References

Bogdan, Radu, 1983, Fodor's Representations, Cognition and Brain Theory, 6, 2, 237-249.

Bogdan, Radu, 1986, The Objects of Perception, in Bogdan (ed), Roderick Chisholm, Dordrecht: Reidel.

Bogdan, Radu, 1988a, Information and Semantic Cognition, Mind and Language, 3, 81-122.

Bogdan, Radu, 1988b, Mental Attitudes and Commonsense Psychology,

⁵ An early and much different version of this paper was read at a Washington University conference on 'Perspectives on the Mind' in December 1991 in response to a paper by Martin Davies. I want to thank Martin and other conference participants, subsequent audiences, and two anonymous referees, for their comments and suggestions.

Nous, 22, 369-398.

Bogdan, Radu, 1991, Common Sense Naturalized, in Bogdan (ed), Mind and Common Sense, Cambridge: Cambridge University Press.

Bogdan, Radu, Guidance to Goal, Hillsdale, NJ.: Lawrence Erlbaum.

Cosmides, L. and Tooby, J. 1989, Evolutionary Psychology and the Generation of Culture, Ethology and Sociobiology, 10, 29-49, 51-97.

Davies, Martin, 1992, Concepts, Connectionism, and the Language of Thought, in D. Rumelhart et al. (eds), Philosophy and Connectionist Theory, Hillsdale, NJ: Lawrence Erlbaum.

Fodor, Jerry, 1975, The Language of Thought, Cambridge: Harvard University Press.

Fodor, Jerry, 1981, Representations, Cambridge, MA: The MIT Press.

Fodor, Jerry, 1983, The Modularity of Mind, Cambridge, MA: The MIT Press.

Fodor, Jerry, 1987, Psychosemantics, Cambridge, MA: MIT Press.

Fodor, Jerry, and Pylyshyn, Zenon, 1988, Connectionism and Cognitive Architecture, Cognition, 28, 3-71.

Pylyshyn, Zenon, 1984, Computation and Cognition, Cambridge, MA: The MIT Press.

Ramsey, William; Stich, Stephen; Garon, Joseph, 1990, Connectionism, Eliminativism and the Future of Folk Psychology, in Philosophical Perspectives, 4, 499-533.

Rumelhart, David, et al., 1986, Schemata and Sequential Thought Processes in PDP Models, in McClelland, J.L. and Rumelhard, D.E., Parallel

Distributed Processes, vol. 2, Cambridge, MA: The MIT Press.

Stich, Stephen, 1983, From Folk Psychology to Cognitive Science, Cambridge, MA: The MIT Press.

Stich, Stephen 1988: From Connectionism to Eliminativism, The Behavioral and Brain Sciences, 11, 53-54.