

MENTAL ATTITUDES AND COMMON SENSE PSYCHOLOGY: THE CASE AGAINST ELIMINATION

RADU J. BOGDAN
TULANE UNIVERSITY

I. INTRODUCTION

Aside from brute force, there are several philosophically respectable ways of eliminating the mental. In recent years the most popular elimination strategy has been directed against our common sense or folk psychological understanding of the mental. The strategy goes by the name of eliminative materialism (or eliminativism, in short). The motivation behind this strategy seems to be the following. If common sense psychology can be construed as the principled theory of the mental, whose vocabulary and principles implicitly define what counts as mental, then eliminating the theory is eliminating its subject matter. If the theory is shown to be false, then its subject matter does not exist. If, in other words, common sense psychology can be shown to describe and explain nothing real in human cognition, then the mental itself is a fiction.

In its earlier stages, in the hands of Rorty and Feyerabend, among others, eliminativism had focused on the common sense notions and accounts of the sensory and phenomenal aspects of cognition. This made tactical sense because it is in the phenomenally accessible arena of consciousness that, for centuries, the dominant Cartesian paradigm has placed the mind. In recent years, however, the critical punch of eliminativism has been redirected against the common sense view of the nonphenomenal, informational and functional states such as belief, thought, desire, and the like. This redirection of critical focus also makes sense, from the standpoint of eliminativism, because it is in the area of cognitive information processing and functions that the scientific action is these

days. Since now it is propositional attitudes, as understood and employed by common sense psychology, that implicitly define the mental (in the nonphenomenal sense), the strategy of eliminativism is to set up a conflict between the sciences of cognition and common sense psychology, which in terms of their respective subject matter becomes a conflict between cognition and mind or, adjectivally, between the cognitive and the mental.

Needless to say, it matters a lot whether eliminativism is right or not about common sense psychology and its subject matter, the mental. Throughout its history, the mental had a way of resurfacing in new clothes after having been exposed and eliminated in the old ones. This is because the mental has firm, possibly permanent functions in our scheme of understanding. It distinguishes and protects me (and any of us) from the surrounding world (subjective/objective); from the physicality of the world, with its flux, change and mortality (soul/body); from you (other minds, anyone?); from animality and "machinity"; from physical and social determination (freedom, creativity, personality, individuality); from being mere information processing and computation (human being/computer). I am enumerating all these familiar facts to underscore what is at stake in the debate about the elimination of common sense psychology, and to suggest the (often hidden) motivation many philosophers would have to come to the rescue of common sense psychology, which they perceive as the last defense of the mental in the face of the ineluctable advance of science, materialism and machine cognition.

This, however, is not going to be my angle on the issue of eliminativism versus common sense psychology. Although programmed to be a common sense psychologist, as we all are, I am not an advocate of its wisdom, nor a defender of its cherished notions. Moreover, contrary to what many philosophers think, being a common sense psychologist does not make one an

expert on common sense psychology. So I do not have a theory of common sense psychology. That must eventually come from social psychology and anthropology, I guess. But I do have three independent working hypotheses, one about common sense psychology, another about what is wrong with eliminativism, and a third about the mental. I want to bring them together. The resulting picture looks as follows.

Eliminativism misrepresents common sense psychology because it misrepresents cognition itself. In particular, eliminativism misrepresents the common sense notions of propositional attitudes, hence the notion of the mental that the latter implicitly define, because it misrepresents the forms in which information is encoded, attitudinized and processed by central cognition. These misrepresentations are strongly encouraged by the axiom that everything about cognition can be explained in a current scientific format. There are reasons to doubt the truth of this axiom. These reasons are objectively motivated by the very ways in which information is shaped and handled in central cognition. I will call this form of information, specific to central cognition, *mental*. The common sense notions of propositional attitudes, when properly reconstructed as attitudes to mental information, can be shown to realistically capture something of importance about central cognition. In brief, the case against eliminativism and for common sense psychology is going to rest on an argument which takes us *from* a novel characterization of mental information and attitudes *to* the inability of the current formats of scientific explanation to account for them and then *to* the realistic legitimacy of the framework which helps us locate and understand mental information and attitudes, namely, common sense psychology. The upshot of the argument is that if central cognition works in ways approximated by the notions of mental information and mental attitudes, then those notions cannot be eliminated. And

if common sense psychology is the principled source of our notions of mental information and attitudes, then it too is bound to escape elimination.

This conclusion seems right not because I take common sense psychology to be the right theory about central cognition or because I take its notions of mental information and attitudes to be the only right notions about central cognition. For I do not. I do not construe common sense psychology as a theory, but rather as a practice. Nor, consequently, do I construe its constitutive notions as being descriptively right about what they implicitly characterize, but rather as practically right. Being practically right about the mind does not deprive the common sense notions of their realistic commitments. Any practice is systematically successful only if the notions that guide it capture real properties of the targets of the practice. To this extent, then, common sense is neither false of cognition, when construed realistically, nor merely heuristic, when construed instrumentalistically, as versions of eliminativism would have it.

This is what I want to argue in what follows. The argument will have some complexity. Having its basic outline from the outset will help in following the different moves and seeing how they serve the overall objective. I first identify the programmatic project of eliminativism and then distinguish several major elimination strategies which carry out this project. These strategies also turn out to be influential accounts of common sense psychology. This is as far as section II takes us. Section III examines the model of cognition, envisaged by all elimination strategies, which satisfies the conceptual and explanatory requirements of the current sciences of cognition. The model, that of transductive and modular (henceforth transmodular) cognition, is the one which eliminativism opposes to the common sense notion of mind. Section IV takes a fresh look at the notion of information, on which everything else rests, and

comes up with a new form, mental information, which is shown to be psychologically real, causally efficacious and constitutive of various mental attitudes. Mental information is also shown to be an irreducible form of information *relative to* the transmodular laws of cognition, hence inherently underdetermined by those laws and thus underexplained by the sciences of transmodular cognition. The concluding section, V, brings common sense psychology back into the picture. The suggestion there is that there is one account of common sense psychology, also contemplated by some eliminativists, the individualist account, which comes closest to construing common sense psychology as an objective, biosocially motivated practice which handles mental information and its associated attitudes. Unlike its eliminativist account, I construe this practice as saying something realistic and important about the human mind. Against eliminativism, the mental is back in the head, and common sense back to a legitimate position, downgraded from theory to practice.

II. A DIAGNOSIS OF ELIMINATIVISM

1. *The Script*. The eliminativist script provides the following basic recipe for eliminating the mental:

(1) ASSUME: Cognition is information computed and represented ¹. The representation and computation of information are objective, law-abiding features of cognizers in interaction with their environment, which is why these features can and should be characterized and explained, in an appropriate format, in the legitimate sciences of cognition.

(2) ASSUME: Common sense characterizes the *mental* in terms of cognitive states such as beliefs, thoughts or desires. Philosophers construe

such states as attitudes to semantic contents and call them *propositional attitudes*. The attitudinal contents and their laws constitute the principled ontology of mental kinds and laws which common sense psychology draws on to conceptualize and explain cognition and behavior.

(3) ASSUME: Only what can be explained *scientifically*, in terms of empirically determined kinds and laws, *is real* at an appropriate level of abstraction. The rest is epiphenomenal.

(4) SHOW: The mental has *nothing* to do with the cognitive. Show this by arguing that the crucial common sense notions of propositional attitudes, which implicitly define the mental, describe nothing *real* in or about cognition.

(5) ALSO SHOW: The sciences of cognition can explain *everything* that is important about cognition, without ever employing any common sense notions such as those of propositional attitudes.

(6) CONCLUDE: Common sense psychology identifies and explains nothing real about cognition. The mental (that common sense psychology implicitly defines) is not scientifically real. The cognitive (that the sciences of cognition implicitly or explicitly define) exclude the mental. Hence the mental does not exist. Elimination accomplished.

2. Stories of elimination. The eliminativist script animates several distinct stories about how science eliminates common sense by purging the cognitive of the mental. I distinguish four stories of elimination, labeled in terms of how they construe common sense psychology and its basic constructs. I identify and discuss these stories only relative to a few parameters which I take to bear on our argument. Detail, completeness and exegetical accuracy are not my objectives here. I summarize the key aspects in a fairly self-explanatory table and then add a few expositive details here and there.

[Insert: Table 1]

The *realist* story is best summarized by its most outspoken author, Paul Churchland. On his view, [realist] eliminativism is

"the thesis that our common sense conception of psychological phenomena constitutes a radically false theory, a theory so fundamentally defective that both the principles and the ontology of that theory will eventually be displaced, rather than smoothly reduced, by completed neuroscience" ².

This entails that the "completed neuroscience" will eventually tell the whole story of human cognition, in which case common sense psychology will turn out to have been just a temporary though sturdy intellectual convenience, the transient byproduct of our ignorance.

Given the parameters chosen to show how eliminativism views common sense, the argument behind the realist version can be reconstructed as follows. Common sense psychology is an *empirical theory* descriptive of human cognition, whence the realist construal. It has its own ontology of kinds and laws which, like in any bona fide theory, provide its main resources of explanation. The *kinds*, the propositional attitudes, are meant to characterize principled types of internal states such as beliefs or thoughts. Common sense psychology contains a rich network of *laws* about propositional attitudes which are neither analytic nor normative. These laws describe objective regularities among attitudes, or between the latter and verbal or behavioral outputs, based on logical relations among the *contents* of the attitudes involved. The common sense *explanations* are of the deductive-nomological sort, as they bring

individual attitudes and behaviors under the relevant laws. Given all this, common sense psychology can be construed as *competing* with the fundamental sciences of cognition, the neurosciences in particular, in individuating and explaining cognitive and behavioral facts. The outcome of the competition is clear. Common sense psychology is an irremediably inadequate theory, outperformed in its own domain by the neurosciences. The neurosciences will therefore *displace* and *replace* it and its explanatory notions. The mind is out, neurocognition is in.

The *interpretational* story also takes common sense psychology to be a *theory*, either anticipative of (Dennett), or continuous with, some appropriate social/normative sciences (Davidson), but *not* an empirical one. Its task is to interpret, rationalize and evaluate cognition in the light of values such as optimality, usefulness or rationality. As a theory, common sense psychology comes equipped with its own ontology of kinds and laws but neither the *kinds* (propositional attitudes which capture normative or teleological aspects of cognition) nor the *laws* (which embody normative or teleological constraints on the relations among contents) are meant to refer to any real, instantiable types of cognitive feature and regularities. Whatever its particular logical structure, a common sense *explanation* is going to be intentional, first, because it treats the mental tokens opaquely, as represented by the agent, and second, because the explanation is designed to make the explanandum appear reasonable or optimal, given the agent's knowledge and interests, and the social norms he conforms to. So there can be no competition between the empirical sciences of cognition and the interpretative common sense psychology, for they are not sharing the same ontology of kinds and laws. Propositional attitudes and their laws fail to have cognitive reality. To this extent, they are *eliminated* from the business of cognition yet remain explanatorily and

predictively important. The mind, as implicitly defined by common sense, is an useful, heuristic though nonreferential explanatory construct. Although no part of cognition, hence unreal, the notion of mind helps us understand and predict cognition.

The *dual-track* version tells a more comprehensive story about what common sense psychology is doing. The internal track monitors what is going on, cognitively, in the head, typically, in the form of functional roles captured by logical relations among contents. This track is deemed insufficient to fix the semantic contents of the attitudes. An external track, which monitors the causal relations with the environment, is also needed to establish the semantic aboutness of contents. What dual-track common sense psychology does is connect mental states with their actual semantic targets, whether in terms of reference, truth conditions or social conventions. The concepts of propositional attitudes capture jointly *kinds* of internal states and their external semantic coordinates. Whatever the *laws* and hence the common sense explanations are, they certainly are not about cognition as such. The notion of the mental is an explanatory/evaluative hybrid which matches cognition at no natural joints. This time the elimination of the mental follows a stipulated reconceptualization of the range of semantic attitudes and contents.

The *individualist* strategy, finally, strikes in a different direction. It is less theoretically motivated than the others yet it surfaces in bits and pieces all over the place³. It construes common sense as a social practice which provides us with tools with which we manufacture individualized case studies of an agent's cognition and behavior. Unlike the other three elimination strategies which portray common sense psychology as some sort of protoscience of an empirical or normative form, dealing with what is general and repeatable in matters cognitive and behavioral, the individualist story portrays common sense as a

sort of protoart or skill which provides imaginative recipes for dealing with what is uniquely individual, novel, different. So construed, common sense psychology is *no theory*, in any reasonable sense. Its notions capture no empirical or normative kinds, indeed *no kinds* at all. The notion of kind specifies some homogeneous class under some distinct property. Read individualistically, propositional attitudes are not at all like that. Their nature is implicitly defined by our ordinary attributions and evaluations in appropriate contexts. If we look carefully at these attributions and evaluations, what we find is a pragmatic potpourri of disjointed and often incommensurate dimensions brought together to fix the contents and the attitudes involved. Since there is no unique nature to be read into those contents and attitudes, it is not very clear what they are, from context to context, nor does this seem to matter. *Laws?* Surely, the individualist would say, common sense helps itself to whatever generalizations serve its purposes. But this does not mean that common sense formulates or explicitly endorses its *own* set of laws. Since the notion of content is so fluid, it would be hard to have any genuine content laws anyway. *Explanation?* Surely, there is explanation, of a sort. But no particular logical format stands out. All that matters is to skillfully pile up enough relevant detail until an individualized profile of a token of one's cognition, motivation or behavior emerges and makes sense relative to whatever norms prevail in the cultural background. There is obviously no competition here with any scientific understanding of cognition. The mental gets *eliminated* nonetheless, for it is now more remote from, and incommensurate with, the cognitive than it ever was.

III. THE ELIMINATIVIST PROFILE OF COGNITION:

DIAGNOSIS CONTINUED

To make its case, eliminativism must not only portray the mental (in the form of propositional attitudes) as epiphenomenal, and the cognitive as real, because the latter but not the former is scientifically characterizable and explainable. It must *also* show that no explanatorily important aspect of cognition is in principle beyond the reach of the sciences of cognition, for if there were such aspects, they could possibly legitimize modes of theorizing and explanation alternative to the scientific ones. Eliminativism must therefore defend the notion that cognition is an object of scientific (specifically, cognitive-theoretical) characterization and explanation *in its entirety*, without any principled residue. Yet the current constraints on explanation in the sciences of cognition are inevitably going to place strong limitations on the profile of cognition envisaged by eliminativism. Can eliminativism have it both ways? Can it bring cognition under science, completely? That is the question we want to answer in this section.

1. Modeling for Explanation: Type, Legislate and Explain. Making scientific sense of a domain of facts comes in two stages. The first is that of conceptualization. It amounts to organizing the facts into kinds (or types) among which pervasive correlations or laws obtain. This means providing the domain with an *ontology* of kinds and laws. The first stage is a prelude to the second, that of explanation, when we locate the token facts we want to understand in the ontology cooked up in the first stage. In so doing, we characterize (or type) those facts according to what the ontology says and then subsume them under appropriate laws. When the logical subsumption is completed, we have explained the facts. If all this is done the Science's way, we have scientifically explained the facts.

To be an object of scientific explanation, cognition must be construed in

ways which allow that sort of explanation to apply to it. This means that the explanation of cognition must pass the scientific test: it must meet some standard requirements of scientific explanation. No exhaustive analysis of such requirements will be attempted here. All we need, for our argument, is to recall only those which reveal the nature of the scientific approach to cognition and point to its limits.

Scientific explanation is explanation *from* kinds and laws: the fact to be explained must bear a satisfactory logical relation to the kinds and laws invoked -- a relation, that is, which satisfies our sense of, and desire for, explanation. The tradition has it that the most respectable format of scientific explanation is *deductive nomological*. It types the explanandum as a particular instance of a kind and (with contextual conditions specified) deduces it from a law. This format seems to work best at lower ontological levels, such as physical and chemical. At more complex levels, where structures embody capacities and dispositions, we need a *morphological* format where we explain capacities in terms of underlying structures and more fundamental capacities. The morphological format seems most apt for complex types of hardware, up to and including the neural ware. When, however, structures and capacities are designed to cooperate in bringing about a result or performing a function, we need a still further format of explanation, the *systematic* ⁴. Not surprisingly, it is the latter that cognitive theoretical explanation needs.

Whatever its format, scientific explanation works only if its domain is modeled, at an appropriate level of abstraction, so as to allow some logical relation between the explanans and the explanandum to be operative. This is done, as we saw, by typing and legislating the right ontology which ensures that the tokens in the domain belong to kinds whose properties and relations are governed by laws bearing on correlations, or primitive capacities, or structures,

or functions, as the case may be. Portraying an explanandum as the regular token of a kind, under some law, ensures that what we know about the kind and the law extends to the token. This is what scientific explanation is all about.

Another requirement is *materialism*. This is a metaphysical requirement which prescribes that all tokens be material (stuff in space-time) at some level of abstraction, and that all kinds at some level have a downward material tokening at some inferior level. All is Ware. In our world, physicalism is the basic version of materialism: all is physical ware.

The next requirement, which I baptize the requirement of *the generic explanandum*, is both delicate and very central to my argument⁵. What it says is that in the scientific explanation of an *individual* event, thing or property, the explanandum is a *generic* individual. What I mean by this is the following. Any material token is going to be *trivially* unique because the aggregation of all its properties make it different from any other token. If *all* the properties of a token are considered, then no two tokens are ever alike, so each is (trivially) unique. But scientific explanation does not care about all the properties of a token. This indifference is deliberate. Trivial unicity cannot and need not be explained. In chasing the universal in the particular, scientific explanation abstracts *precisely* from those properties which render a token trivially unique.

Any token explanandum can be both trivially unique and generic. It is trivially unique when all its properties are taken into account; it is generic when what makes it trivially unique is ignored. A token is generic when, for a given context of space-time-causation, any token of the same kind, occupying the same explanandum position, gets the same explanation. For scientific explanation, the context is really unique, not its occupier. The context identifies its occupier for the explanans and in so doing it confers trivial unicity upon it. The context itself is not explained. It is the occupier which is explained but only

as a generic item, as a faceless representative of its kind. As far as scientific theorizing is concerned, any similarly sized and constituted particle or molecule or crystal or is going to behave the same way as any other, no matter what space-time-cause context it occupies. They are generic individuals because the theory models them as faceless, unoriginal, average, law-abiding members of uniform, homogeneous kinds governed by laws. The point, you would notice, is not about what is out there, but about how theories model, in order to explain, what is out there.

The notion of the generic explanandum is to be later contrasted with that of a nongeneric or *essentially unique* explanandum. The latter notion is about targets of explanation which are essentially unique either because the very contexts they occupy make a lot of explanatory difference or because the explananda themselves have properties and dispositions responsible for nonlawful, uniquely creative and original behavior. Typically, the two go hand in hand: contexts matter explanatorily when their occupiers display nonlawful and original reactions to causal inputs. This distinguishes trivial from essential unicity: space-time- causal contexts confer trivial unicity on a token explanandum when its behavior remains invariant under changes of context; when the token's behavior does not remain thusly invariant, it shows sensitivity to context, and if such sensitivity matters explanatorily, we have an essentially unique explanandum.

2. Cognitive Theoretical Explanation . Cognitive theoretical explanation is a species of scientific explanation. It is the scientific explanation of cognition, at appropriate levels of ontological abstraction, in formats suitable to those levels. The levels themselves, as defined by the sciences of cognition, such as psychophysics, neuroscience or cognitive psychology, specify the domains of

cognitive theoretical explanation. In addition to the general requirements discussed a section ago, which make it *scientific*, cognitive theoretical explanation must also satisfy a number of domain specific requirements, which make it *cognitive theoretical*.

A first requirement affects the very modeling of the ontology needed by cognitive theoretical explanation. It concerns specifically the cognitive kinds recruited for conceptualization and explanation in the sciences of cognition. The requirement is known as *methodological solipsism*. It stipulates that we type cognitive kinds, such as representations, beliefs or intentions, without regard to the actual aspects of the environment they may be (when tokened) semantically about. This does not exclude taking into account the systematic influences of the environment, society or evolution (as the requirement is often and mistakenly construed). Those influences had better be considered if we want to understand what cognition is doing in this world and how. The requirement only excludes *current* and *particular* semantical coordinates from counting in the type individuation of cognitive kinds.

There is also the requirement of *internalism*. It urges that any *informational* transaction in which an organism is involved should either explain or be explained in a cognitive theoretical manner only to the extent to which the transaction is tokened internally by states of the organism. The motivation here is simple. To explain cognitive and behavioral tokens is to explain what causally drives cognition and behavior. Information is that driving force. An information token is causally efficacious in an organism only if, as long as, and in the very form in which, it is encoded by the states of the organism.

3. *What and How Much Gets Explained?* Let us now put the requirements together and see what sort, and how much, of cognition they allow cognitive

theoretical explanation to handle. This in turn should tell us how much of our understanding of the cognitive mind will ultimately originate in the sciences of cognition. I will ask the question in the following form: What should human cognition be like if it were to meet the requirements on cognitive theoretical explanation? Where could the explanation recruit its kinds and laws from in order to bring them to bear on its explananda?

Let us begin with the second question. The answer must point to an organism's hardware and its basic semantic design: they alone deliver the kinds and laws that the sciences of cognition need, at different levels, to do their explanatory job. The basic semantic design of cognition is essentially embodied in transducers, (representation-generating) modules, control and motor centers. These are the components which enable the organism to organize and process the input in ways which direct its behavior toward aspects of the environment which are of interest. The information tokened in the organism can then be said to be *about* those aspects of interest, hence to be semantic. This is dogmatically short but then so is the space allotted to this paper. The longer story is told in [3].

Now the answer to the first question. If human cognition were just transmodular cognition, run on a suitable hardware and sufficient to engage the control and motor centers and guide behavior toward its assigned types of targets, then cognitive theoretical explanation would apply fully and successfully to *any* instance of cognition and behavior. If, in other words, cognition were just *semantic cognition* (in the naturalized sense anticipated a paragraph ago), then the sciences of cognition would alone be in the business of explaining it. Or look at it this way. Suppose all that matters in understanding cognition is how information is converted (or transduced) and encoded (or represented) in order to provide a semantic map for behavior, in the sense that

the information tells the control and motor centers what types of targets are available out there to be acted on or reacted to. The information that does this job does a semantic job. That, moreover, is the sort of job that (I suggest) the current sciences of cognition are perfectly capable of conceptualizing and explaining. What I am saying is that transmodular cognition is essentially semantic cognition, for it lawfully delivers semantic representations of the environment, and as such is intelligible to, and explainable by, the sciences of cognition.

Why is semantic cognition a neat object of cognitive theoretical explanation? Because any token of it can be thusly explained. But why is this so? Because any semantic information token (when typed in compliance with the requirements on cognitive theoretical explanation) falls under kinds of hardware and/or semantic design which portray it as generic in the explanandum position and obeying the right laws. We are all alike in our transmodular cognition. We are built that way. If we find the right kinds and laws behind the way we are built, we are in the cognitive theoretical business. This is precisely what people have been doing for the last few decades in the areas of perception, language, learning, memory and so on. What their efforts have in common is understanding how semantic information is transduced, encoded and operated on so as to service the semantic objectives of the organism (i.e. what its behavior must, by design, be sensitive to). Fodor's story [10] of transmodular cognition tells it all and much better than I can tell it here.

4. Interim Review. What does all this amount to? If transmodular cognition is the only object of cognitive-theoretical explanation, then, for the eliminativist, it must be the theory of transmodular cognition that (implicitly) defines the *cognitive* and excludes the *mental*. The transmodular is not the

mental or, differently put, the mind cannot be transmodular. This is a sensible conclusion. So far eliminativism is right. But now recall the other half of its project: to show that the *entirety* of cognition can be explained cognitive-theoretically. Yet what can be so explained is just transmodular cognition. Eliminativists do know that there are other sorts of cognitive states and processes such as those involved in thinking, problem solving, deliberating or speech acting which are not transmodular. They are central. The choice for eliminativism is either to belittle the significance of central cognition, or to argue that central cognition could be characterized and explained (away) in a bottom-up fashion, in terms of transmodular cognition, plus some additional hypotheses. Most serious eliminativists, whatever their story of mind is, would probably choose the second option. If this bottom-up project were successful, then the cognitive-theoretical account could be extended to central cognition ⁶.

The test is the notion of information. Cognition is an information driven enterprise. It is the notion of information which is at the center of the conflict between eliminativism and common sense psychology. The eliminativist project can succeed only if the information operative in central cognition can be characterized and explained in transmodular terms. Can it? The next two sections attempt to answer this question.

IV. ANOTHER LOOK AT INFORMATION

1. "*Propositional Attitudes*". The case for eliminativism and against common sense psychology rests on the pivotal notion of propositional attitude. The notion of proposition is semantic, for it is meant to capture the semantic (or aboutness) aspects of the information handled in cognition. The semantic properties of information are real and important, but why should we assume

that they constitute the very ultimate form of information which drives our cognition and behavior? The information tokened in cognition is neither exclusively semantic nor fully operative just in virtue of its being semantic. That information also has physical properties, and syntactic properties, and other properties as well. Each type of property is tokened at some level. Yet neither the physical nor the syntactic nor the semantic level is necessarily the (final) level where attitudes are formed and behaviors initiated. A formal computation in the brain is physical but we know that a physical account will not tell us what that computation is and what it is doing in the brain. The same can be said about a semantic encoding of information and its syntactic structure: an account of the latter need reveal neither the nature nor the cognitive function of the former. The widespread semantic fallacy in contemporary philosophy of mind, backed by no good argument and lacking any explanatory value, is to hold that the semantic form of information (i.e. information encoded under semantic and intentional constraints) enters *as semantic* into functional positions and initiates behavior. This is very wrong.

Before seeing why it is so wrong, let us pause to note how the semantic fallacy helps eliminativism. If cognitive information were fully analyzable as semantic information, then either its habitat would be transmodular cognition (because this is where semantic information is assembled) or else, wherever semantic information gets constituted, it would be within the reach of transmodular constraints. In either case, the processes involved would be those which the sciences of cognition can fully understand and explain. The plot, then, is clear: if one buys the semantic fallacy, so successfully embodied in the notion of propositional attitude, one is already buying a good deal of eliminativism.

So what exactly is a propositional attitude? The notion, attributed to common sense psychology, is in fact a philosophical concoction. Common sense

works with the more specific notions of belief, desire, intention, among many others. Philosophers abstract from these specific notions a common form which, since Russell, is construed as an attitude to a semantic content, typically analyzed as proposition. Even when analyzed differently, say, as a sentence or syntactic formula or some other sort of representation, the analytic intention is still that of capturing the *semantic* aspects of the information encoded in a representation. It is this semantic intention of the philosophical analysis that must be challenged. The questions we must ask are: Is the semantic notion of propositional attitude a fair explication of the common sense notions? And does the semantic notion of propositional attitude make independent explanatory sense? To answer these questions without begging them we need a few distinctions.

To begin with, let us be neutral and call the notion we are after *information attitude*. There are two major components in an information attitude, the attitude and the information. Each can be characterized at different levels of abstraction. Since information is the critical notion in our discussion I will bracket out questions about attitudes. I assume that we are designed to have the attitudes we have and that they interact in equally well designed ways. Even if we have a full scientific map of the "mechanics" of attitude formation and interaction, I take it that it is still the objects of those attitudes (the information) that matter most to common sense psychology ⁷.

So back to information. There are two sides to information in cognition, one internal, the other external. The intuitive distinction is this. Information is a structure or an encoding produced in an organism by a proximal input originating in an external source. An aspect of information is *internal* (hence cognitive) only if it is encoded, in some form, in a state of the organism's hardware and it is the encoding which is causally efficacious in cognition. An

aspect of information is *external* if it is productive of cognition (as the light input is) but is not internal. Now I call *semantic* those external aspects or conditions of information such as reference, truth and input conditions, etc., which characterize the fact as well as the degree of interaction between internal encodings of information (say, representations) and their external objects. I take both such interactions and various social rules, conventions, practices and contextual factors to be constraints on semantic information. Finally, I construe concepts, meanings, prototypes, recognition patterns and the like as *intentional* aspects of information, internal to cognition. The intentional aspects are either explicitly represented in memory, as data structures, or are implicitly followed, as acquired or built-in rules and procedures. The important point is that, although they act as constraints on semantic information (in the sense that they capture semantic facts and relations), the intentional dimensions operate on information cognitively, internally. A semantic information token, then, is *jointly* typed by both intentional and external constraints and facts. It is this joint contribution that determines what the token is semantically about.

The standard semantic notion of propositional attitude is that, when an information token is attitudinized, it is its *semantic* properties that are attitudinized and become functionally efficacious. When (impressed by the case made in the *New Yorker*) I come to believe that, having been invited to join the faculty at Princeton, George Washington was wondering how long it takes to get tenure there, the standard view wants my belief to be sensitive to and engage only the semantic aspects of the information tokened in my head on this occasion. The standard view holds, in other words, that the information which is the object of my belief can be fully explicated by an aggregation of external and intentional parameters which fix its semantic aboutness. The

standard view also holds that the paramount relations among, and laws of, propositional attitudes are semantic in nature because they are sensitive to truth preservation (a semantic virtue). The key content relations among beliefs are thought to be logical, hence semantically sensitive, in the form of implication, equivalence, and so forth. The laws governing the relations among beliefs have the normative and again semantically sensitive form of (say) consistency, coherence or closure. And so on. Such is the standard view, so dominant and so wrong.

2. Readings of Information . Eliminativism skillfully exploits the semantic reading of information attitudes by first attributing it to common sense and then showing it to be, on the one hand, explanatorily useless, if construed purely externally, and, on the other hand, subject to a cognitive theoretical explanation, if construed purely internally. The first hand is easy to play. The external reading of semantic information cannot by itself explain anything cognitive and behavioral. So we must wonder why common sense, after thousands of years of daily experimentation and feedback, and a fairly successful record of attitude attribution and rationalization, would assume otherwise. The requirement of internalism, which common sense psychology appears to follow, is violated. This is because no external aspect of information by itself has any internal causal efficacy; and when it has, because encoded internally, it is no longer just semantic. Semantic notions such as reference, truth conditions and other indexicalities are used to *evaluate* how internal encodings of information fare when compared with their external targets. Since common sense does constantly engage in semantic evaluation, it makes more sense to assume that its use of semantic notions is *evaluational* rather than descriptive and explanatory of cognition. I will come back to this point.

The other hand reads information attitudes internally. There are at least three major internal readings of information: in terms of hardware, syntactic and intentional constraints. The hardware reading is too weak to type individuate cognitive information, for it fails to specify its formal encoding and functional role. Both encoding and role are in principle compatible with many types of ware. The argument is very familiar, so we need not pause on it. We also notice that common sense doesn't much care for ware, either. So we move on. The syntactic reading gives the nuts and bolts of the internal representation of information. Many people, including some eliminativists (e.g. Stich [13]), think that the syntactic reading is the *final* internal reading of information. This is the syntactic fallacy. It entails that there is no further level of abstraction at which we can fruitfully (that is, with an eye to explanation) type individuate internal properties of information. The syntactic fallacy serve the eliminativist cause even better than the semantic fallacy because the syntax of cognitive representation, unlike its semantics, fully complies with internalism and methodological solipsism and hence is a legitimate object of cognitive-theoretical explanation. Common sense has nothing illuminating to say about it and indeed does not care much about it. If cognitive information is just a matter of hardware and syntax, common sense has no privileged conceptual hold on information. So it can be abandoned.

Yet the syntactic reading cannot by itself fix cognitive information, either. A superficial and familiar reason is that the same syntactic structure (*The nuts are all over the campus*) may encode different semantic contents, just as different syntactic structures (say, in different languages) may encode the same semantic content. Even a finer-grained sublinguistic syntax is likely to face the same problem, as we will see in a moment. The more serious reason is that cognitive syntax must encode information in formats which, while respecting

semantic aspects of interest to the organism, are conducive to behavior. This semantic and functional business of the syntax cannot be explained in syntactic terms, without adverting to its motivating factors.

To see what else is needed, let us ask a counterfactual question: What would be the conditions in which the syntactic fallacy could turn into a true thesis about the mind? These would be the conditions in which the syntactic reading counts as the final internal reading of information which fixes the form of information which drives an organism's cognition and behavior. What sort of organism would make the syntactic fallacy true? Well, the computer is always there to help. Think of one, named HAM, which never knows what it is doing. HAM is a computer whose syntactic types today track (say) flight connections and tomorrow check the airline's personnel names, weights and subscriptions to libertarian publications. The semantic interpretation of what HAM is doing, today as well as tomorrow, is ours. All that HAM does is compute. It computes but does not represent anything. It *cannot* represent. Or so we say.

What is the implied difference between computation and representation? What is it about the information that HAM tokens that makes it syntactic but not semantic? Think of it this way. The reason HAM can work on flight connections is that it has been programmed in such a way that its syntactic types and operations encode and preserve relevant semantic aspects in the intended domain of interpretation (flight connections). This syntactic-semantic correlation or alignment is not something that HAM knows. HAM was not told about the correlation, nor is it endowed with the internal resources to token or learn the correlation by aligning its syntactic forms and operations to the semantic aspects we care about. It is the internal tokening (in some form) of the alignment of syntax to semantics that makes all the difference in the world of cognition, a tokening which HAM is deprived of. This is what makes HAM a

mere syntactic engine.

Or put it this way. HAM has the syntactic ability to compute flight connections but it does not have the *further* ability to type individuate flight connections *as such*, and distinguish them from, say, personnel files, or from anything else, for that matter. Poor HAM does not have *concepts*. This is why, in turn, the syntactic forms which HAM computes represent nothing in particular and therefore have no *meaning*. A syntactic form can be said to have meaning only if it is deployed under some concepts or other classificatory routines designed (by learning, evolution, God, what have you) to map syntactic patterns onto semantic aspects. To read a content-expressing syntax in the light of concepts and meanings is to read it *intentionally*. Concepts can be viewed (in this context) as functions from syntactic forms to semantic aspects, as constraints which ensure, first, that the syntactic encoding of information is sensitive to the semantic aspects of significance to the cognizer and, second, that the computation of information preserves those semantic aspects. Concepts have this task because they must ensure that the organism's behavior is directed toward specific environmental targets. This is awfully sketchy but I hope it will do, particularly if it can help answer our next question.

How does all this bear on our story about eliminativism, common sense and information attitudes? The intentional aspects of cognition, concepts and meanings in particular, are not and need not be the proprietary business of common sense psychology. Common sense does not seem to worry about or inquire into the very nature or function of the basic concepts and meanings we normally employ -- unless something goes wrong. This suggests that common sense psychology is not a genuine theory of our basic intentional program. But the intentional program *can* be explained cognitive-theoretically. And so it is. Psychological theories of learning, for example, study concept formation and

application. Theories of vision and language study how the syntax engages the semantic aspects which matter most in seeing and saying things. And so on.

Where science takes hold, common sense takes leave. If common sense psychology has no principled interest in our intentional design, and if the latter can be explained by the sciences of cognition, we ought to wisely infer that common sense cannot durably and ineliminably be in the intentional business. And if common sense persists in having a psychological business of its own, we ought to wisely infer that it must lie *elsewhere*. Granted, these are not inferences easy to make, even by those philosophers who defend common sense, for they too think that if common sense illuminates and exploits something about the mind, it is precisely its intentionality. (The semantic fallacy is the bread and butter of contemporary analytic philosophy.) And I am denying that, for I am arguing that intentionality (or mental aboutness) is barely the beginning of the common sense story of the mind.

If common sense is not primarily in the intentional business, what business is it in, if it is to deal with cognitive information in internal terms? If the information tokened by a central state cannot be fixed, separately or jointly, by either its externally semantic conditions or its hardware or syntactic or intentional parameters, then what else is there in the head to help with the fixing?

3. *Mental Information.* There is indeed more to cognitive information than current philosophical wisdom acknowledges. Or so I want to suggest. The long story is told elsewhere ([1], [2] and [3]). We get here only a serviceably short abstract of it. I have earlier introduced the new form of information as *mental* information. This form has an internal encoding but of a pragmatic nature. It is *this* form of encoding, NOT its underlying syntactic and semantic features,

which is functionally efficacious in cognition and behavior. This is why the semantic and syntactic fallacies misrepresent cognition at work.

To see what this means, let us begin by distinguishing the notion of mental information from other important and familiar notions of information. We must first distinguish it from the notion of *input information* which Fred Dretske has recently analyzed so well [9]. The latter is about the interactions between external sources and organisms, specifically, about the information that physical signals deliver at the sensory gates. This is a *physical* form of information encoded as hardware patterns. The input information does specify how the organism represents the input and how that representation produces further representations and actions. For that we need the *semantic* form of information. Sensitive to syntactic and intentional constraints, this form is defined over semantic aspects of interest to the organism, and specifies how the latter aspects fit the former constraints. In particular, the semantic information encoded by a representation can specify which states of affairs are compatible with that representation, and which are not. This brings the notion of semantic information very close to that of meaning, and makes it characterizable up to the intentional level of abstraction.

We saw that this is not going to be enough to understand cognition. We must also determine how the semantic possibilities made available by the input information and identified by its representational encoding are going to affect, say, what a cognizer's uncertainty about a situation is, what she already knows about it, how the uncertainty is reduced, and how she acts on the new information. Change this uncertainty scenario to solving a problem, making a decision, choosing a hypothesis, communicating something, or acting -- the basic story remains the same. It is, in all these cases, a story in which the semantic information made available by input and encoded under syntactic and

intentional constraints *is relativized*, under new constraints, to the current cognitive state of the cognizer. The result is mental information.

Take communication, for an example. I hear an utterance. Its meaning allows me to represent its semantic possibilities, or truth conditions. But I do not believe, and act on, meanings alone. I also expect an utterance to make a *further* contribution to my cognition by interacting with what I know, ignore, expect, and plan about a situation. It is in this *further* interaction, which goes beyond semantics and intentionality, that mental information is shaped; and it is shaped pragmatically and incrementally, according to the aggregate cognitive state in which I happen to be and the behavioral objectives I pursue.

The facts about mental information are embarrassingly but significantly familiar. You and I may be in a situation in which, upon receiving (type) similar inputs, visual or verbal, we react differently. Such cognitive and behavioral differences may exist even when key internal aspects of information fixed by those (type) similar inputs also happen to be (type) similar. Imagine, for example, that both you and I hear the malicious rumor, so worded, that Robert Vesco has offered to rebuild the Tulane basketball program. We may plausibly assume that, phonetically, gramatically and conceptually, we both process that input statement in type similar ways. Assuming that our concepts of Vesco, rebuilding, basketball, and program are alike, we read the same literal meaning into that statement. Yet, after computing that meaning, you say, "Aha, that is interesting", whereas I see red stars and broken teeth. Not quite the same inferences, evaluations and behavior.

Why this difference? I happen to know more about the Tulane basketball program than you do. I may also have different expectations and fears about this basketball issue than you have. The difference in expectations and fears may help explaining the difference in behavior. But there must be *antecedently*

a difference in information which explains a number of cognitive moves (inferences drawn, beliefs fixed) I make and you don't. If so, the information you and I extract from the input statement is going to be different. This is mental information.

Imagine our HAM endowed with perceptual analyzers, formal rules and meaning-fixing concepts. HAM can access some input information p about a situation s , can represent p syntactically in a formula f , and can bring f under suitable concepts which fix f 's meaning, m . In this story HAM registers, computes and represents some input information up to the intentional level. So far HAM can mimic the perceptual, syntactic and intentional processes we go through when understanding a sentence. Suppose now that in HAM's case, unlike ours, the intentional output $f(m)$ goes nowhere, for HAM has no memory and no behavior that its memory could serve. As soon as it computes and represents some $f(m)$, that $f(m)$ just vanishes. It interacts with no prior knowledge, no beliefs, no expectations, no plans, nothing. HAM's current cognitive state makes no difference to what happens to the intentional output $f(m)$.

Our human $f(m)$, however, is further subject to an interaction with our current cognitive states (beliefs, expectations, and the like). Since information is a matter of structures created in interaction with other structures, this *further* type of interaction, governed by a *new* kind of constraints (action-oriented and pragmatic), should alert us to the existence of a new form of information. Lacking the resources for such an internal interaction, HAM is unable to fix the mental information that $f(m)$ would have added to HAM's current configuration of stored beliefs. But there is no such configuration in HAM, so no information is added to it and no belief fixed. Occasionally, each of us comes close to being a HAM when, out of the blue, we register an isolated

input, say, an utterance which we can figure out up to the intentional (or meaning) level, i.e. as an $f(m)$, but no further, because we cannot bring our knowledge or expectations to bear on the meaning of the utterance. The notion of mental information is meant to capture precisely the difference between what is represented, as a meaning or $f(m)$, and what that representation adds to what we already know, expect and anticipate.

Why is this important? Why does mental information matter? Information, we saw, always comes in the form of a structure generated by interactions with other structures under appropriate laws and constraints. The reason semantic information (or content) is type different from mental information is that the former is generated by input under hardware, syntactic and intentional constraints, whereas the latter is generated by the semantic information in further interaction with current cognitive and behavioral states under new sorts of constraints (of an inductive and pragmatic nature, pertaining to what is stored, believed, planned, expected, etc.) Mental information has the pragmatic job of aligning semantic information to the organism's current cognition and action. The semantic information, jointly fixed by causal interactions (with the environment) and by syntactic and intentional constraints, cannot alone do this job. Put a bit metaphorically, the semantic information looks back at the environment to represent it, whereas mental information takes the representation and looks forward at how it can service current and future cognition and action. Semantic information is thus constrained by world-to-mind relations, mental information by mind-to-action-to-world relations.

The fixation of mental information is an inductive process in which various data-encoding states interact with a target representation (the interim intentional output) and fix a specific increment in the cognizer's knowledge, given what he already knows, expects, and wants to do. The fixation process

must be holistic, adhoc and contextual, hence highly individualized, because so are the contributing interactions. Given their job, the constraints on these mental-inductive processes promise no invariance, the way syntactic and intentional constraints do. The processes which fix mental information could not be invariant because they are underdetermined not only by context and input but also by the internal constraints at the hardware, syntactic and intentional levels. The fixation of mental information is a genuinely central and pragmatic process ⁸.

V. MENTAL ATTITUDES BRIEFLY REVISITED

In order to bring our story of mental information to bear on our understanding of common sense psychology we must first relate it to our understanding of mental attitudes. The semantic fallacy has it that mental attitudes have semantic contents as their objects. This does not make much sense. Mental attitudes must be geared to how information animates and moves one's current cognition and behavior. And, we saw, it isn't semantic information *as semantic* which does the animating and the moving. We believe and desire because we have to act; and in order to act we must token the information which aligns the action to our current cognitive state. In other words, we treat information believingly and desiringly because that information must service current cognition and action.

This is a brute teleological fact. HAM does not have beliefs and desires because it is under no obligation to act; but it can process information up to the semantic level. True, for very good reasons, we often *treat* beliefs and other attitudes as (backward-looking) relations to semantic contents only, ignoring

their (forward-looking) relations to current and future cognition and action. But that is just an opportunistic tactic, not evidence of what beliefs and other attitudes really are. I can treat a book as a hammer but that does not make a book a hammer. I just exploit one property of books to convince nails to enter the wall, while ignoring many other properties books have more essentially. The semantic dimensions of a belief, the semantic information or content it encodes, may be dominantly important in some context, say, of truth evaluation or intentional interpretation. In those cases, a belief may be *treated as* a relation to a semantic content (say, a proposition). For important biosocial reasons, common sense is sensitive to semantic evaluation and intentional interpretation and often treats beliefs accordingly. But it is not in the nature of a belief to be just semantic, or else its cognitive business is totally obscured.

The issue may seem merely terminological. We may call 'belief' the (merely) intentional attitude to a semantic content and something else the attitude to mental information. This terminological manoeuvre should not however assume that the cognitive mind tokens and operates with *both* sorts of attitudes, or else we end up with two unnecessary coordination problems. One concerns the coordination between the semantic/intentional half and the mental half of "belief", the other the coordination between whatever now counts as semantic/intentional belief and other kinds of attitudes such as desires and intentions. (The latter is a problem because desires and intentions are geared to action whereas intentional belief is not. When intentional belief is action-oriented, it is no longer just intentional, it is mental.)

In real life cognition both coordination problems have a simple and natural solution because only the tokening of the mental attitudes is real and causally efficacious. The intentional attitude to a semantic content, by itself, has no cognitive reality and efficacy; it is merely the practically motivated artifact of

attribution (on which more later) or the philosophically motivated artifact of phenomenological and metalinguistic analysis. In his natural cognitive stance, a cognizer can hardly attitudinize a semantic content just intentionally (or under concepts), without at the same time putting it to further cognitive and behavioral work. Once put to such work the content is no longer just semantic, and the attitude no longer just intentional, since their very structure and operation are now regulated by further, nonsemantic and nonintentional constraints.

To summarize, my strategy has been, first, to let mental facts and regularities constrain semantic information, just as we let semantic and intentional facts and regularities constrain underlying syntactic encodings and operations; and second, to declare cognitively real and causally efficacious the type of information and attitudes which get typed and constrained last. To be as successful as it is in handling the pragmatics of cognition and behavior, common sense must acknowledge and exploit this very truth about mental information and attitudes. This is what remains to be argued in the conclusion of our argument.

VI. COMMON SENSE PSYCHOPRAXIS

If common sense makes sense of an agent's cognition and behavior in terms of attitudinized mental information, then the strategies employed by common sense to this effect must be individualistic yet realistic because so is the very nature of attitudinized mental information, namely, cognitively real and causally efficacious yet constituted in essentially individualized and pragmatic conditions. I think it takes no fancy argument to show, first, that common sense is interested (among other things) in explaining individual cognition and behavior

in terms of attitudinized mental information; and second, that common sense is capable of producing individualist stories about how particular agents fix information, to which they extend their cognitive and noncognitive attitudes and on which they then act. The fixation of mental information is essentially (as opposed to trivially) individual and pragmatic because it is relative to the aggregate cognitive and behavioral state of an agent, in the form of rather ephemeral increments. The shape of the information is incremental because, as we saw, input always meets knowledge stored in memory, with the result that only the novel increment that the former adds to the latter becomes cognitively relevant and efficacious. Yet the increment is recognized by the cognizer only in a particular thematic context, in a given conceptual format, relative to other possible increments compatible with what the cognizer knew, was uncertain about and expected. The incrementation is ephemeral because all its parameters, just listed, change values from context to context.

Let us now put the resulting picture together. We have genuine and important facts in cognition in the form of mental information under various attitudes. Attitudinized mental information drives cognition and behavior in essentially individualized and pragmatic ways which must be understood if cognition and behavior are to be understood. We also have an intriguingly pervasive and successful paradigm of conceptualization and explanation of cognition and behavior, which we call common sense psychology, whose explanatory interests seem to go in the direction of attitudinized mental information and whose explanatory strategies seem to fit the object of its interest. We can recall at this point that there is a plausible methodological assumption, discussed in section III, to the effect that the ontological nature of the object of theorizing very much dictates how it ends up conceptualized and explained. Given all these elements, it is natural to entertain the hypothesis that

the individualist strategies of common sense psychology are calibrated to fit the nature of its essentially individual and pragmatic object of explanation. If this hypothesis is plausible, then eliminativism must be wrong about at least one form of the mental and its conceptual home, common sense psychology.

Yet the argument, plausible as it may seem so far, will not carry much conviction if we do not attempt to make a bit more sense of common sense psychology itself. This is our next and last task. Recall, to begin with, an earlier observation (section IV.3) to the effect that the facts about mental information are significantly familiar. The reason why this is so has to do with the very nature of common sense psychology. If we find it so easy and familiar to track and specify (attitudinized) mental information, it is because it is second nature to us. What common sense enables us to do so well, effortlessly and unreflectively, is a matter of routine or practice, of skill, rather than of theory. It is a practice which recognizes mental information in its various attitudinal embodiments without necessarily recognizing or caring about its syntactic, semantic or even intentional underpinnings. This may be no different from our practical knowledge of, say, tables and chairs or making omlette, which is opaque to and uninterested in the deep constitution of its objects.

If we come to understand that common sense psychology is an implicit, unreflective practice directed at a family of mental phenomena, rather than an explicit theory about those phenomena, let alone about their deeper, necessary conditions, then we can also see why 'common sense psychology' is quite a misnomer. There is no *logos* in it, only *praxis*. Common sense is psychopraxis, not psychologos. It is a practice which, for its successful exercise, relies on concepts and conceptual correlations which capture various objective properties and regularities of interest. To handle a car I need a number of concepts (speed, acceleration, steering, braking, etc.) and a number of interconceptual

correlations about various regularities (turning the ignition in the parking position starts the engine (though not in my old car), braking on wet surfaces is slippery, and so on), both at some rather superficial level of description. Both my handling and my knowledge of cars is practical, not theoretical. I do not study cars, I just do things with/to them, based on some minimal practical knowledge. That knowledge has realistic import, for it is about cars and the situations in which they operate. In an even more implicit manner, the same is true of my practice of riding a bike, swimming or indeed walking. Similarly, cats could not catch mice if their hunting practices were not guided by realistic notions about mice behavior. Those notions need not form a theory of mice behavior to do their job. So with common sense psychopraxis: to do its job successfully, it has to capture explanatorily real, causally efficacious aspects of central cognition. If a practice had no realistic bite, what would be its point?

Yet tracking the fixation of mental information, figuring out the attitudes toward it and hence the cognitive and behavioral outputs it generates seems to be just one of the jobs of common sense. Another job, when the need arises, is intentional interpretation, that is, the attempt to figure out the conceptual repertory of another cognizer, from the same tribe or another. Translation and radical interpretation, as understood by philosophers since Quine, are dramatic examples of intentional interpretation. The job here is equally psychological but it concerns a form of cognitive competence, our conceptual program. By contrast, the psychopraxis which handles attitudinized mental information is concerned mostly with how the conceptual program is deployed in actual cases of cognition, as a matter of cognitive performance. Still another line of business for common sense is, as we saw, the semantic evaluation of mental states, relative to their natural and social environments. This, however, is no longer a psychological job. And then there is epistemic evaluation. All these are distinct

practices of common sense which focus on separate areas of our cognitive life. The practices reflect a sensible division of labor motivated by biosocial reasons. Their story must await another occasion. How we come to learn, master and exercise these practices is still a mystery. We switch easily from the psychopractical stance of tracing one's fixation of mental information to the semantic stance of identifying the truth conditions of that information or to the intentional stance of ascertaining what concepts are being used. And so on.

Philosophers who take common sense and ordinary language to implicitly define various aspects of our mental life often make the mistake of taking a particular common sense practice as speaking for common sense in general. They take that specific practice as embodying the general common sense understanding of the mental. For many reasons, semantic evaluation and intentional interpretation are special favorites. In philosophy of mind and language, they get to define, for example, what mental attitudes are -- which is why we end up with the semantic and/or intentional notion of propositional attitude wrongly paraded as fully explicative of the cognitive mind.

The argument of this paper, on behalf of mental attitudes and common sense, and against eliminativism, has focused only on one type of common sense practice, a psychological practice concerned with mental information in attitudinized roles. It is *this* practice which, against eliminativism, vindicates a realist but individualist and pragmatic reading of mental attitudes to information. On this reading, it appears that the fixation of mental attitudes and their causal efficacy are real facts which must but cannot be satisfactorily handled by the current formats of cognitive theoretical explanation. Yet it does not follow that there is competition between common sense and the sciences of cognition for our ultimate understanding of mental affairs. Either the sciences deliver such understanding or they don't because they cannot. In the latter

case, common sense cannot step in to take up the slack. What common sense does is identify the phenomena it must figure out, for reasons of its own, even if they are to remain unexplained scientifically. Only when a philosophical position like eliminativism claims that common sense psychology is false about the mind because there are no such things as mental attitudes, and does so in the name of the sciences of cognition, only then do we have a conflict between science and common sense. But the conflict is not about explanation, it is about the ontology of cognition.

Common sense psychology is a very potent practice with no apparent programmatic interest in the inner workings of cognition. If my overall argument is on the right track, this superficiality is not unmotivated. Knowledge of the lawful inner workings of cognition, up to the intentional level, would not help identify and explain mental attitudes. We have shown this to be true, in principle, *independently* of the story of common sense. Yet it seems that common sense implicitly recognizes that central cognition has information fixing capabilities which, on the one hand, are underdetermined by internal mechanisms and laws of the hardware, formal and intentional sort (whence its superficiality); but which, on the other hand, are quite sensitive to a lawless, pragmatic variety of data, both in the head and outside it (whence its practical, contextual, multi-track, improvisational approach to mental information). If the performances generated by those capabilities, in the form of mental attitudes and acts, are what an individualist but realist common sense psychology implicitly defines and explains from, then the mental is part of the cognitive. The mental specifies performance states of central cognition in the form of thoughts, beliefs, intentions and inferences.

The mental, then, is back in the head, up there where the sciences of cognition seem unable to reach it. This is not because common sense is

descriptively right but rather because objectively there is something up there in the head anyway, something which biosocial life has forced common sense to acknowledge and handle practically. This is true no matter what we think of common sense. The intriguing thought, however, is that in some implicit, unreflective, practical way common sense may have known all along that the hardware, syntax and semantics of cognition do not tell the whole story⁹.

REFERENCES

- Bogdan, Radu J., 1987, "Mind, Content, and Information", *Synthese*, 70: 205-227.
- _____ (ed), 1986, *Belief*, Oxford: Oxford University Press.
- _____, 1986, "The Manufacture of Belief", in [Bogdan 1986].
- _____, 1985, "The Intentional Stance Reexamined", *The Behavioral and Brain Sciences*, 8: 759-760.
- _____, *Mind and Information*, forthcoming.
- Churchland, Paul, 1979, *Scientific Materialism and the Plasticity of Mind*, Cambridge: Cambridge University Press.
- _____, 1981, "Eliminative Materialism and Propositional Attitudes", *Journal of Philosophy*, 78: 67-90.
- Churchland, Paul and Patricia, 1983, "Stalking the Wild Epistemic Engine", *Nous*, 17: 6-18.
- Dennett, Daniel, 1978, *Brainstorms*, Cambridge, MA: MIT Press.
- Dretske, Fred, 1981, *Knowledge and the Flow of Information*, Cambridge, MA: The MIT Press.
- Fodor, Jerry, 1983, *The Modularity of Mind*, Cambridge, MA: The MIT Press.

- Haugeland, John, 1978, "The Nature and Plausibility of Cognitivism", *The Behavioral and Brain Sciences*, 2: 215-225.
- Morton, Adam, 1980, *Frames of Mind*, Oxford: Oxford University Press.
- Stich, Stephen S., 1983, *From Common Sense Psychology to Cognitive Science*, Cambridge, MA: The MIT Press.
- Wilkes, K.V., 1981, "Functionalism, Psychology, and the Philosophy of Mind", *Philosophical Topics*, 12: 147-168.

NOTES

¹ The nature of the representations, whether analog or digital, symbolic or nonsymbolic, would not concern us here, although a decision on this issue could affect the way we would ultimately explain cognition, in terms of what sciences of cognition, at what level of abstraction.

² [Churchland 1981, p. 67]. All I say about the realist version, including the synopsis of its basic argument to be developed in the next paragraph, is based on Paul Churchland's works, particularly [Churchland 1979 and 1981]. Yet in both my synopsis and the discussion following it there is an element of generalization which goes beyond Churchland's specific neuroscientific eliminativism. This is because one can be a realist eliminativist at a level of abstraction higher than that of the neurosciences, say, at the level of cognitive psychology.

³ [Adam Morton's 1980] and papers by K.V. Wilkes, particularly her [1981], articulate in detail an individualist version. Some of the Wittgensteinian and Austinian ordinary language literature of the '50's and 60's also embodies eliminativist prejudices of the individualist sort about common sense.

Eliminativist views may share several of the versions of common sense psychology I have distinguished here. Many do. Wittgenstein's own, for example, may be construed as a mixture of the interpretational and individualist versions.

4 The distinction between these formats of scientific explanation is Haugeland's in [Haugeland 1978].

5 In correspondence, Paul Churchland has indicated that he takes both the typical scientific and the common sense psychological explanation to be concerned with explaining unique individual events and properties. I agree but draw (in a moment) a distinction between trivial unicity and essential unicity. Churchland considers only the former. He attributes it to common sense psychology because he takes common sense psychology to be a protoscience.

6 Paul Churchland and perhaps Dan Dennett as well seem to think that a theory of transmodular cognition plus considerations about evolution and its impact on our cognitive design will do the cognitive-theoretical bottom-up explanatory trick [Churchland&Churchland 1983; Dennett 1978]. The basic idea seems to be this. The design of central cognition is responsible for the properties of mental information and its fixation. So the explanation of the former will explain the latter as well. How, then, could we explain the design of central cognition? Think of how we explain the transmodular design. The latter, according to Churchland, can be treated as "something that evolution calibrates (i.e., as something which, by random mutation and on random selection, is tuned to measure, via the excitable cells, certain features in the environment)". For the central design we must assume that evolution selects, not registration

patterns and associated motor responses, but rather learning strategies as such. What we are assuming, in other words, is that "the organism has become so 'well-tuned' that it seems to have *an evolving world-picture*, rather as though the organism has tricked up an analogue of the evolutionary process itself". Learning, then, is evolution internalized. Given this evolutionary assumption about central design, Churchland concludes: "Here then the job for neurobiology and neuroethology is Herculean but the bets are that the story for complex organisms will be build on the more basic story of calibrational [read: modular] semantics for simpler organisms, following the steps of evolution itself" [Churchland&Churchland 1983, p.14] .

My response is to grant Churchland the assumption (that evolution can explain central design) but not the conclusion (that central design can explain mental performances). So let us suppose that Churchland is right about evolution "calibrating" the central design the ways he says it does. There are two questions we want to ask about this supposition. First, *how far up* along the levels of abstraction at which we must examine cognition would the supposition about the calibrating evolution take us? And second, once at the highest level we can reach, *how much* can we explain from central design? The most optimistic answer to the first question is: as far up as the intentional level, where concepts and meanings work. If human cognition were totally regularized by (or "closed under") intentional properties and laws, then indeed tokens of cognition could be fully characterized at that level and, given some initial conditions, could also be explained just from central design. But our story of mental information will show that this is not so. The intentional competence which the central design specifies *under* determines the fixation and the functional consequences of mental information. Explanation of central competence is not yet explanation of its mental performances. This in turn

answers the second question. I further develop this response in [Bogdan, 1985] in a discussion of Dennett's views.

⁷ There are plenty of attitudinal laws which surface in our common sense psychological explanations. I am thinking of such laws as *If S fears that p, then S desires that not p*; or *If S hopes that p and finds out that p, then S is pleased that p*; and so on. Some such laws involve only attitudes, other attitudes and actions. Among eliminativists, [Churchland 1979, and 1981] and [Morton 1980] discuss attitudinal laws in great detail. I call such laws *attitudinal*, not content, laws because the source of their lawfulness is the attitudes, not the content involved. It is a matter of attitudinal design that fear goes with desire, whatever their contents. A content law, on the other hand, even though governing information attitudes, is lawful because of the contents those attitudes have. For example, *If S believes that p, and also believes that p implies q, then S believes that q* is an informational law because it is based on the logical law of modus ponens which the belief attitude tracks, for whatever reason. In an attitudinal law it is the other way around: the content *p* in a law to the effect that *hope that p + finding out that p = pleasure that p* tracks an attitudinal regularity. I deliberately talk of content, as a semantic notion, not of information, to emphasize the prevalent semantic prejudice in the analysis of cognitive attitudes.

Most attitudinal laws are specified by our basic design, while some are perhaps learned. Common sense is aware of and exploits such laws. Common sense may even be wrong about the attitudinal laws. Yet whether it is right or not about them, common sense psychology does not seem, *pace* Churchland, to be a theory of attitudinal laws, no more than it is a theory of grammar or vision or concept learning. If anything, common sense psychology is more likely

to be concerned with information. My point, therefore, is to concede to Churchland, Dennett, and others that the "mechanics" of attitude interaction and its laws could become the subject matter of a science of cognitive design. I would also concede that common sense psychology may be all wrong about that "mechanics". But I still insist that common sense psychology is primarily interested in attitudinized *information*, not in the attitudes themselves. When it comes to information, however, the cognitive design is much less helpful, as the previous note shows.

⁸ I develop this in much more detail in [Bogdan, forthcoming].

⁹ I have started a predecessor to this paper some time ago. It was meant to be a discussion and refutation of the realist version of eliminativism, particularly Paul Churchland's. Then, stimulated by further thought plus a number of observations and criticisms, I came to consider the larger picture of eliminativism versus common sense psychology. Some of the initial points have survived, others have not. A number of people have seen or heard some draft or another, some development or another, at some stage or another, and have made critical comments and suggestions which helped a lot. I want to thank, very much, Paul Churchland for his careful reading of an earlier draft and the detailed comments he made on it; an anonymous referee for *Nous* who made some very good points; friends and colleagues at the CNRS-sponsored seminars in Paris and at the University of Helsinki who have heard and reacted to papers echoing things said here; and, as always, my friends at Tulane and the University of New Orleans, Harvey Green, Carolyn Morillo and Norton Nelkin in particular, who, after hearing installments of this story on several occasions, may think that I am funded by Hoi Polloi, Inc. If I were, I would be in Paris.

