The Philosophical Review, XCV, No. 1 (January 1986)

INDIVIDUALISM AND PSYCHOLOGY*

Tyler Burge

Recent years have seen in psychology—and overlapping parts of linguistics, artificial intelligence, and the social sciences—the development of some semblance of agreement about an approach to the empirical study of human activity and ability. The approach is broadly mentalistic in that it involves the attribution of states, processes and events that are intentional, in the sense of 'representational'. Many of these events and states are unconscious and inaccessible to mere reflection. Computer jargon is prominent in labeling them. But they bear comparison to thoughts, wants, memories, perceptions, plans, mental sets and the like—ordinarily so-called. Like ordinary propositional attitudes, some are described by means of that-clauses and may be evaluated as true or false. All are involved in a system by means of which a person knows, represents, and utilizes information about his or her surroundings.

In the first part of this paper, I shall criticize some arguments that have been given for thinking that explanation in psychology is, and ought to be, purely "individualistic." In the second part of the paper, I shall discuss in some detail a powerful psychological theory that is not individualistic. The point of this latter discussion will be to illustrate a non-individualistic conception of explanatory kinds. In a third section, I shall offer a general argument against individualism, that centers on visual perception. What I have to say, throughout the paper, will bear on all parts of psychology that attribute intentional states. But I will make special reference to explanation in cognitive psychology.

Individualism is a view about how kinds are correctly individuated, how their natures are fixed. We shall be concerned primarily with individualism about the individuation of mental kinds. According to individualism about the mind, the mental natures of all

*A version of this paper was given at the Sloan Conference at MIT in May 1984. I have benefited from the commentaries by Ned Block, Fred Dretske, and Stephen Stich. I have also made use of discussion with Jerry Fodor, David Israel, Bernie Kobes, and Neil Stillings; and I am grateful to the editors for several suggestions.

a person's or animal's mental states (and events) are such that there is no necessary or deep individuative relation between the individual's being in states of those kinds and the nature of the individual's physical or social environments.

This view owes its prominence to Descartes. It was embraced by Locke, Leibniz, and Hume. And it has recently found a home in the phenomenological tradition and in the doctrines of twentieth century behaviorists, functionalists, and mind-brain identity theorists. There are various more specific versions of the doctrine. A number of fundamental issues in traditional philosophy are shaped by them. In this paper, however, I shall concentrate on versions of the doctrine that have been prominent in recent philosophy of psychology.

Current individualistic views of intentional mental states and events have tended to take one of two forms. One form maintains that an individual's being in any given intentional state (or being the subject of such an event) can be *explicated* by reference to states and events of the individual that are specifiable without using intentional vocabulary and without presupposing anything about the individual subject's social or physical environments. The explication is supposed to specify—in non-intentional terms—stimulations, behavior, and internal physical or functional states of the individual. The other form of individualism is implied by the first, but is weaker. It does not attempt to explicate anything. It simply makes a claim of *supervenience:* an individual's intentional states and events (types and tokens) could not be different from what they are, given the individual's physical, chemical, neural, or functional histories, where these histories are specified non-intentionally and in a way that is independent of physical or social conditions outside the individual's body.

In other papers I have argued that both forms of individualism are mistaken. A person's intentional states and events could (counterfactually) vary, even as the individual's physical, functional (and perhaps phenomenological) history, specified non-intentionally and individualistically, is held constant. I have offered several arguments for this conclusion. Appreciating the strength of these arguments, and discerning the philosophical potential of a non-individualist view of mind, depend heavily on reflecting on differences among these arguments. They both reinforce one another and help map the topography of a positive position.

For present purposes, however, I shall merely sketch a couple of the arguments to give their flavor. I shall not defend them or enter a variety of relevant qualifications. Consider a person *A* who thinks that aluminum is a light metal used in sailboat masts, and a person *B* who believes that he or she has arthritis in the thigh. We assume that *A* and *B* can pick out instances of aluminum and arthritis (respectively) and know many familiar general facts about aluminum and arthritis. *A* is, however, ignorant of aluminum's chemical structure and micro-properties. *B* is ignorant of the fact that arthritis cannot occur outside of joints. Now we can imagine counterfactual cases in which *A* and *B*'s bodies have their same histories considered in isolation of their physical environments, but in which there are significant environmental differences from the actual situation. *A*'s counterfactual environment lacks aluminum and has in its places a similar-looking light metal. *B*'s counterfactual environment is such that no one has ever isolated arthritis as a specific disease, or syndrome of diseases. In these cases, *A* would lack "aluminum thoughts" and *B* would lack "arthritis thoughts." Assuming natural developmental patterns, both would have different thoughts. Thus these differences from the actual situation show up not only in the protagonist's relations to their environments, but also in their intentional mental states and events, ordinarily so-called. The arguments bring out variations in obliquely (or intensionally) occurring expressions in literal mental state and event ascriptions, our primary means of identifying intentional mental states.[1]

I believe that these arguments use literal descriptions of mental

[1] "Individualism and the Mental," *Midwest Studies* 4 (1979), pp. 73–121; "Other Bodies," in *Thought and Object,* Woodfield, ed. (Oxford: Oxford University Press, 1982); "Two Thought Experiments Reviewed," *Notre Dame Journal of Formal Logic* 23 (1982), pp. 284–293; "Cartesian Error and the Objectivity of Perception," forthcoming in *Subject, Thought, and Context,* MacDowell and Pettit eds. (Oxford: Oxford University Press, 1986); "Intellectual Norms and Foundations of Mind" (forthcoming, *The Journal of Philosophy*). The aluminum argument is adapted from an argument in Hilary Putnam, "The Meaning of 'Meaning'," *Philosophical Papers* Vol. II (Cambridge, England: Cambridge University Press, 1975). What Putnam wrote in his paper was, strictly, not even compatible with this argument. (Cf. the first two cited papers in this note for discussion.) But the aluminum argument lies close to the surface of the argument he does give. The arthritis argument raises rather different issues, despite its parallel methodology.

events, and are independent of conversational devices that may affect the form of an ascription without bearing on the nature of the mental event described. The sort of argument that we have illustrated does not depend on special features of the notions of arthritis or aluminum. Such arguments go through for observational and theoretical notions, for percepts as well as concepts, for natural-kind and non-natural kind notions, for notions that are the special preserve of experts, and for what are known in the psychological literature as "basic categories." Indeed, I think that, at a minimum, relevantly similar arguments can be shown to go through with any notion that applies to public types of objects, properties, or events that are typically known by empirical means.[2]

I shall not elaborate or defend the arguments here. In what follows, I shall presuppose that they are cogent. For our purposes, it will be enough if one bears firmly in mind their conclusion: mental states and events may in principle vary with variations in the environment, even as an individual's physical (functional, phenomenological) history, specified non-intentionally and individualistically, remains constant.

A common reaction to these conclusions, often unsupported by argument, has been to concede their force, but to try to limit their effect. It is frequently held that they apply to common-sense attributions of attitudes, but have no application to analogous attributions in psychology. Non-individualistic aspects of mentalistic attribution have been held to be uncongenial with the purposes and requirements of psychological theory. Of course, there is a tradition of holding that ordinary intentional attributions are incapable of yielding any knowledge at all. Others have held the more modest view that mentalistic attributions are capable of yielding only knowledge that could not in principle be systematized in a theory.

I shall not be able to discuss all of these lines of thought. In particular I shall ignore generalized arguments that mentalistic ascriptions are deeply indeterminate, or otherwise incapable of

[2]On basic categories, cf., e.g., Rosch, E., Mervis, Gray, Johnson, Boyes-Graem, "Basic Objects in Natural Categories," *Cognitive Psychology* 8 (1976), pp. 382–439. On the general claim in the last sentence, cf. "Intellectual Norms," *op. cit.* and the latter portion of this paper.

yielding knowledge. Our focus will be on arguments that purport to show that non-individualistic mentalistic ascriptions cannot play a systematic role in psychological explanation—*because* of the fact that they are not individualistic.

There are indeed significant differences between theoretical discourse in psychology and the mentalistic discourse of common sense. The most obvious one is that the language of theoretical psychology requires refinements on ordinary discourse. It not only requires greater system and rigor, and a raft of unconscious states and events that are not ordinarily attributed (though they are, I think, ordinarily allowed for). It also must distill out descriptive-explanatory purposes of common attributions from uses that serve communication at the expense of description and explanation. Making this distinction is already common practice. Refinement for scientific purposes must, however, be systematic and meticulous—though it need not eliminate all vagueness. I think that there are no sound reasons to believe that such refinement cannot be effected through the development of psychological theory, or that effecting it will fundamentally change the nature of ordinary mentalistic attributions.

Differences between scientific and ordinary discourse survive even when ordinary discourse undergoes the refinements just mentioned. Although common sense discourse—both about macrophysical objects and about mental events—yields knowledge, I believe that the principles governing justification for such discourse differ from those that are invoked in systematic scientific theorizing. So there is, *prima facie,* room for the view that psychology is or should be fully individualistic—even though ordinary descriptions of mental states are not. Nevertheless, the arguments for this view that have been offered do not seem to me cogent. Nor do I find the view independently persuasive.

Before considering such arguments, I must articulate some further background assumptions, this time about psychology itself. I shall be taking those parts of psychology that utilize mentalistic and information-processing discourse pretty much as they are. I assume that they employ standard scientific methodology, that they have produced interesting empirical results, and that they contain more than a smattering of genuine theory. I shall not prejudge what sort of science psychology is, or how it relates to the natural

sciences. I do, however, assume that its cognitive claims and, more especially, its methods and presuppositions are to be taken seriously as the best we now have in this area of inquiry. I believe that there are no good reasons for thinking that the methods or findings of this body of work are radically misguided.

I shall not be assuming that psychology *must* continue to maintain touch with common sense discourse. I believe that such touch will almost surely be maintained. But I think that empirical disciplines must find their own way according to standards that they set for themselves. Quasi-*apriori* strictures laid down by philosophers count for little. So our reflections concern psychology as it is, not as it will be or must be.

In taking psychology as it is, I am assuming that it seeks to refine, deepen, generalize and systematize some of the statements of informed common sense about people's mental activity. It accepts, for example, that people see physical objects with certain shapes, textures, and hues, and in certain spatial relations, under certain specified conditions. And it attempts to explain in more depth what people do when they see such things, and how their doing it is done. Psychology accepts that people remember events and truths, that they categorize objects, that they draw inferences, that they act on beliefs and preferences. And it attempts to find deep regularities in these activities, to specify mechanisms that underly them, and to provide systematic accounts of how these activities relate to one another. In describing and, at least partly, in explaining these activities and abilities, psychology makes use of interpreted that-clauses and other intensional constructions—or what we might loosely call "intentional content."[3] I have seen no sound reason to believe that this use is merely heuristic, instrumentalistic, or second class in any other sense.

I assume that intentional content has internal structure—something like grammatical or logical structure—and that the parts of

[3]Our talk of intentional "content" will be ontologically colorless. It can be converted to talk about how that-clauses (or their components) are interpreted and differentiated—taken as equivalent or non-equivalent—for the cognitive purposes of psychology. Not all intentional states or structures that are attributed in psychology are explicitly propositional. My views in this paper apply to intentional states generally.

this structure are individuated finely enough to correspond to certain individual abilities, procedures, or perspectives. Since various abilities, procedures, or perspectives may be associated with any given event, object, property, or relation, intentional content must be individuated more finely than the entities in the world with which the individual interacts. We must allow different ways (even, I think, different primitive ways) for the individual to conceive of, or represent any given entity. This assumption about the fine-grainedness of content in psychology will play no explicit role in what follows. I note it here to indicate that my skepticism about individualism as an interpretation of psychology does not stem from a conception of content about which it is already clear that it does not play a dominant role in psychology.[4]

Finally, I shall assume that individualism is *prima facie wrong* about psychology, including cognitive psychology. Since the relevant parts of psychology frequently use attributions of intentional states that are subject to our thought experiments, the language actually used in psychology is not purely individualistic. That is, the generalizations with counterfactual force that appear in psychological theories, given their standard interpretations, are not all individualistic. For ordinary understanding of the truth conditions, or individuation conditions, of the relevant attributions suffices to verify the thought experiments. Moreover, there is at present no well-explained, well-understood, much less well-tested, individualistic language—or individualistic reinterpretation of the linguistic forms currently in use in psychology—that could serve as surrogate.

Thus individualism as applied to psychology must be revisionistic. It must be revisionistic at least about the language of psycho-

[4]Certain approaches to intensional logic featuring either "direct reference" or some analogy between the attitudes and necessity have urged that this practice of fine-structuring attitudinal content be revised. I think that for purely philosophical reasons these approaches cannot account for the attitudes. For example, they do little to illumine the numerous variations on Frege's "paradox of identity." They seem to have even less to recommend them as prescriptions for the language of psychology. Some defenses of individualism have taken these approaches to propositional content to constitute the opposition to individualism. I think that these approaches are not serious contenders as accounts of propositional attitudes and thus should be left out of the discussion.

logical theory. I shall be developing the view that it is also revisionistic, without good reason, about the underlying presuppositions of the science. To justify itself, individualism must fulfill two tasks. It must show that the language of psychology should be revised by demonstrating that the presuppositions of the science are or should be *purely* individualistic. And it must explain a new individualistic language (attributing what is sometimes called "narrow content") that captures genuine theoretical commitments of the science.

These tasks are independent. If the second were accomplished, but the first remained unaccomplishable, individualism would be wrong; but it would have engendered a new level of explanation. For reasons I will mention later, I am skeptical about such wholesale supplementation of current theory. But psychology is not a monolith. Different explanatory tasks and types of explanation coexist within it. In questioning the view that psychology is individualistic, I am not *thereby* doubting whether there are some subparts of psychology that conform to the strictures of individualism. I am doubting whether all of psychology as it is currently practiced is or should be individualistic. Thus I shall concentrate on attempts to fulfill the first of the two tasks that face someone bent on revising psychology along individualistic lines. So much for preliminaries.

I.

We begin by discussing a general argument against non-individualistic accounts. It goes as follows. The behavior of the physiologically and functionally identical protagonists in our thought experiments is identical. But psychology is the science (only) of behavior. Since the behavior of the protagonists is the same, a science of behavior should give the *same* explanations and descriptions of the two cases (by some Ockhamesque principle of parsimony). So there is no room in the discipline for explaining their behavior in terms of different mental states.[5]

[5]Stephen Stich, *From Folk Psychology to Cognitive Science* (Cambridge, Mass.: MIT Press, 1983), chapter VIII. Although I shall not discuss the unformulated Ockhamesque principle, I am skeptical of it. Apart from

The two initial premises are problematic. To begin with the first: it is not to be assumed that the protagonists are behaviorally identical in the thought experiments. I believe that the only clear, general interpretation of 'behavior' that is available and that would verify the first premise is 'bodily motion'. But this construal has almost no relevance to psychology as it is actually practiced. 'Behavior' has become a catch-all term in psychology for observable activity on whose description and character psychologists can reach quick "pretheoretical" agreement. Apart from methodological bias, it is just not true that all descriptions that would count as "behavioral" in cognitive (social, developmental) psychology would apply to both the protagonists. Much behavior is intentional action; many action specifications are non-individualistic. Thought experiments relevantly similar to those which we have already developed will apply to them.

For example, much "behavioral" evidence in psychology is drawn from what people say or how they answer questions. Subjects' utterances (and the questions asked them) must be taken to be interpreted in order to be of any use in the experiments; and it is often assumed that theories may be checked by experiments carried out in different languages. Since the protagonists' sayings in the thought experiments are different, even in non-transparent or oblique occurrences, it is *prima facie* mistaken to count the protagonists "behaviorally" identical. Many attributions of non-verbal behavior are also intentional and non-individualistic, or even relational: she picked up the apple, pointed to the square block, tracked the moving ball, smiled at the familiar face, took the money instead of the risk. These attributions can be elaborated to produce non-individualist thought experiments. The general point is that many relevant specifications of behavior in psychology are intentional, or relational, or both. The thought experiments indicate that these specifications ground non-individualist mental attributions. An argument for individualism cannot reasonably *assume* that these specifications are individualistic or ought to be.

---

question-begging assumptions, it seems to me quite unclear why a science should be required to explain two instances of the same phenomenon in the same way, particularly if the surrounding conditions that led to the instances differ.

Of course, there are non-individualistic specifications of behavior that are unsuitable for any scientific enterprise ('my friend's favorite bodily movement'). But most of these do not even appear to occur in psychology. The problem of providing reasonable specifications of behavior cannot be solved from an armchair. Sanitizing the notion of behavior to meet some antecedently held methodological principle is an old game, never won. One must look at what psychology actually takes as "behavioral" evidence. It is the responsibility of the argument to show that non-individualistic notions have no place in psychology. Insofar as the argument assumes that intentional, non-individualistic specifications of behavior are illegitimate, it either ignores obvious aspects of psychological practice or begs the question at issue.

The second step of the argument also limps. One cannot assume without serious discussion that psychology is correctly characterized as a science (only) of behavior. This is, of course, particularly so if behavior is construed in a restrictive way. But even disregarding how behavior is construed, the premise is doubtful. One reason is that it is hardly to be assumed that a putative science is to be characterized in terms of its evidence as opposed to its subject matter. Of course, the subject matter is to some extent under dispute. But cognitive psychology appears to be about certain molar abilities and activities some of which are propositional attitudes. Since the propositional attitudes attributed do not seem to be fully individuable in individualistic terms, we need a direct argument that cognitive psychology is not a science of what it appears to be a science of.

A second reason for doubting the premise is that psychology seems to be partly about relations between people, or animals, and their environment. It is hard to see how to provide a natural description of a theory of vision, for example, as a science of behavior. The point of the theory is to figure out how people do what they obviously succeed in doing—how they see objects in their environment. We are trying to explain relations between a subject and a physical world that we take ourselves to know something about. Theories of memory, of certain sorts of learning, of linguistic understanding, of belief formation, of categorization, do the same. It is certainly not obvious that these references to relations between subject and environment are somehow inessential to

(all parts of) psychological theory. They seem, in fact, to be a large part of the point of such theory. In my view, these relations help motivate non-individualistic principles of individuation (cf. Section II). In sum, I think that the argument we have so far considered begs significant questions at almost every step.

There is a kindred argument worth considering: the determinants of behavior supervene on states of the brain. (If one is a materialist, one might take this to be a triviality: "brain states supervene on brain states.") So if propositional attitudes are to be treated as among the determinants of behavior, they must be taken to supervene on brain states. The alternative is to take propositional attitudes as behaviorally irrelevant.[6]

This argument can, I think, be turned on its head. Since propositional attitudes are among the determinants of our "behavior" (where this expression is as open-ended as ever), and since propositional attitudes do not supervene on our brain states, not all determinants of our "behavior" supervene on our brain states. I want to make three points against the original argument, two metaphysical and one epistemic or methodological. Metaphysics first.

The ontological stakes that ride on the supervenience doctrine are far less substantial than one might think. It is simply not a "trivial consequence" of materialism about mental states and events that the determinants of our behavior supervene on the states of our brains. This is because what supervenes on what has at least as much to do with how the relevant entities are individuated as with what they are made of. If a mental event m is individuated partly by reference to normal conditions outside a person's body, then, regardless of whether m has material composition, m might vary even as the body remains the same.

---

[6]I have not been able to find a fully explicit statement of this argument in published work. It seems to inform some passages of Jerry Fodor's "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology" in Fodor's *Representations* (Cambridge, Mass.: MIT Press, 1981), e.g., pp. 228–232. It lies closer to the surface in much work influenced by Fodor's paper. Cf., e.g., Colin McGinn, "The Structure of Content" in Woodfield ed. *Thought and Object* (Oxford: Clarendon Press, 1982), pp. 207–216. Many who like McGinn concede the force of the arguments against individualism utilize something like this argument to maintain that individualistic "aspects" of intentional states are all that are relevant to psychological explanation.

Since intentional phenomena form such a large special case, it is probably misleading to seek analogies from other domains to illustrate the point. To loosen up the imagination, however, consider the Battle of Hastings. Suppose that we preserve every human body, every piece of turf, every weapon, every physical structure and all the physical interactions among them, from the first confrontation to the last death or withdrawal on the day of the battle. Suppose that, counterfactually, we imagine all these physical events and props placed in California (perhaps at the same time in 1066). Suppose that the physical activity is artifically induced by brilliant scientists transported to earth by Martian film producers. The distal causes of the battle have nothing to do with the causes of the Battle of Hastings. I think it plausible (and certainly coherent) to say that in such circumstances, not the Battle of Hastings, but only a physical facsimile would have taken place. I think that even if the location in Hastings were maintained, sufficiently different counterfactual causal antecedents would suffice to vary the identity of the battle. The battle is individuated partly in terms of its causes. Though the battle does not supervene on its physical constituents, we have little hesitation about counting it a physical event.

Our individuation of historical battles is probably wrapped up with intentional states of the participants. The point can also be made by reference to cases that are clearly independent of intentional considerations. Consider the emergence of North America from the ocean. Suppose that we delimit what count as constituent (say, micro-) physical events of this larger event. It seems that if the surrounding physical conditions and laws are artfully enough contrived, we can counterfactually conceive these same constituent events (or the constituent physical objects' undergoing physically identical changes in the same places) in such a way that they are embedded in a much larger land mass, so that the physical constituents of North America do not make up any salient part of this larger mass. The emergence of North America would not have occurred in such a case, even though its "constituent" physical events were, in isolation, physically identical with the actual events. We individuate the emergence of continents or other land masses in such a way that they are not supervenient on their physical constituents. But such events are nonetheless physical.

In fact, I think that materialism does not provide reasonable restrictions on theories of the role of mentalistic attributions in psychology. The relation of physical composition presently plays no significant role in any established scientific theory of mental events, or of their relations to brain events. The restrictions that physiological considerations place on psychological theorizing, though substantial, are weaker than those of any of the articulated materialisms, even the weak compositional variety I am alluding to. My point is just that rejecting individualistic supervenience does not entail rejecting a materialistic standpoint. So materialism *per se* does nothing to support individualism.[7]

The second "metaphysical" point concerns causation. The argument we are considering in effect simply assumes that propositional attitudes (type and token) supervene on physico-chemical events in the body. But many philosophers appear to think that this assumption is rendered obvious by bland observations about the etiology of mental events and behavior. It is plausible that events in the external world causally affect the mental events of a subject only by affecting the subject's bodily surfaces; and that nothing (not excluding mental events) causally affects behavior except by affecting (causing or being a causal antecedent of causes of) local states of the subject's body. One might reason that in the anti-individualistic thought experiments these principles are violated

[7]In "Individualism and the Mental," *op. cit.*, pp. 109–113, I argue that token *identity* theories are rendered implausible by the non-individualistic thought experiments. But token identity theories are not the last bastion for materialist defense policy. Composition is what is crucial.

It is coherent, but I think mistaken, to hold that propositional-attitude attributions non-rigidly pick out physical events: so the propositional attributions vary between the actual and counterfactual protagonists in the thought experiments, though the ontology of mental event tokens remains identical. This view is compatible with most of my opposition to individualism. But I think that there is no good reason to believe the very implausible thesis that mental events are not individuated ("essentially" or "basically") in terms of the relevant propositional-attitude attributions. (cf. ibid.) So I reject the view that the same mental events (types or tokens) are picked out under different descriptions in the thought experiments. These considerations stand behind my recommending, to the convinced materialist, composition rather than identity as a paradigm. (I remain unconvinced.)

insofar as events in the environment are alleged to differentially "affect" a person's mental events and behavior without differentially "affecting" his or her body: only if mental events (and states) supervene on the individual's body can the causal principles be maintained.

The reasoning is confused. The confusion is abetted by careless use of the term 'affect', conflating causation with individuation. Variations in the environment that do not vary the impacts that causally "affect" the subject's body may "affect" the individuation of the information that the subject is receiving, of the intentional processes he or she is undergoing, or of the way the subject is acting. It does not follow that the environment causally affects the subject in any way that circumvents its having effects on the subject's body.

Once the conflation is avoided, it becomes clear that there is no simple argument from the causal principles just enunciated to individualism. The example from geology provides a useful countermodel. It shows that one can accept the causal principles and thereby experience no bewilderment whatsoever in rejecting individualism. A continent moves and is moved by local impacts from rocks, waves, molecules. Yet we can conceive of holding constant the continent's peripheral impacts and chemically constituent events and objects, without holding identical the continent or certain of its macro-changes—because the continent's spatial relations to other land masses affect the way we individuate it. Or take an example from biology. Let us accept the plausible principle that nothing causally affects breathing except as it causally affects local states of the lungs. It does not follow, and indeed is not true, that we individuate lungs and the various sub-events of respiration in such a way as to treat those objects and events as supervenient on the chemically described objects and events that compose them. If the same chemical process (same from the surfaces of the lungs inside, and back to the surfaces) were embedded in a different sort of body and had an entirely different function (say, digestive, immunological, or regulatory), we would not be dealing with the same biological states and events. Local causation does not make more plausible local individuation, or individualistic supervenience.

The intended analogy to mental events should be evident. We

may agree that a person's mental events and behavior are causally affected by the person's environment only through local causal effects on the person's body. Without the slightest conceptual discomfort we may individuate mental events so as to allow distinct events (types or tokens) with indistinguishable chemistries, or even physiologies, for the subject's body. Information from and about the environment is transmitted only through proximal stimulations, but the information is individuated partly by reference to the nature of normal distal stimuli. Causation is local. Individuation may presuppose facts about the specific nature of a subject's environment.

Where intentional psychological explanation is itself causal, it may well presuppose that the causal transactions to which its generalizations apply bear some necessary relation to some underlying physical transactions (or other). Without a set of physical transactions, none of the intentional transactions would transpire. But it does not follow that the kinds invoked in explaining causal interactions among intentional states (or between physical states and intentional states—for example, in vision or in action) supervene on the underlying physiological transactions. The same physical transactions in a given person may in principle mediate, or underly, transactions involving different intentional states—if the environmental features that enter into the individuation of the intentional states and that are critical in the explanatory generalizations that invoke those states vary in appropriate ways.

Let us turn to our epistemic point. The view that propositional attitudes help determine behavior is well entrenched in common judgments and in the explanatory practices of psychology. Our arguments that a subject's propositional attitudes are not fixed purely by his or her brain states are based on widely shared judgments regarding *particular* cases that in relevant respects bring out familiar elements in our actual psychological and common sense practices of attitude attribution. By contrast, the claim that none of an individual's propositional attitudes (or determinants of his behavior) could have been different unless some of his brain states were different is a metaphysical conjecture. It is a modal generalization that is not grounded in judgments about particular cases, or (so far) in careful interpretation of the actual explanatory and

descriptive practices of psychology. Metaphysical ideology should either conform to and illuminate intellectual praxis, or produce strong reasons for revising it.

What we know about supervenience must be derived, partly, from what we know about individuation. What we know about individuation is derived from reflecting on explanations and descriptions of going cognitive practices. Individuative methods are bound up with the explanatory and descriptive needs of such practices. Thus justified judgments about what supervenes on what are *derivative* from reflection on the nature of explanation and description in psychological discourse and common attitude attributions. I think that such judgments cannot be reasonably invoked to restrict such discourse. It seems to me therefore that, apart from further argument, the individualistic supervenience thesis provides no reason for requiring (pan-) individualism in psychology. In fact, the argument from individualistic supervenience begs the question. It *presupposes* rather than establishes that *individuation—hence explanation and description*—in psychology should be fully individualistic. It is simply the wrong sort of consideration to invoke in a dispute about explanation and description.

This remark is, I think, quite general. Not just questions of supervenience, but questions of ontology, reduction, and causation generally, are epistemically posterior to questions about the success of explanatory and descriptive practices.[8] One cannot reasonably criticize a purported explanatory or descriptive practice primarily by appeal to some prior conception of what a "good entity" is, or of what individuation or reference should be like, or of what the overall structure of science (or knowledge) should turn out to look like. Questions of what exists, how things are individuated, and

[8]The points about ontology and reference go back to Frege, *Foundations of Arthmetic*, Austin trans. (Northwestern University Press, Evanston, 1968). The point about reduction is relatively obvious, though a few philosophers have urged conceptions of the unity of science in a relatively aprioristic spirit. At least as applied to ontology, the point is also basic to Quine's pragmatism. There are, however, strands in Quine's work and in the work of most of his followers that seem to me to let a preoccupation with physicalism get in the way of the Fregean (and Quinean) pragmatic insight. It is simply an illusion to think that metaphysical or even epistemic preconceptions provide a standard for judging the ontologies or explanatory efforts of particular sciences, deductive or inductive.

what reduces to what, are questions that arise by reference to going explanatory and descriptive practices. By themselves, proposed answers to these questions cannot be used to criticize an otherwise successful mode of explanation and description.[9]

Of course, one might purport to base the individualist supervenience principle on what we know about good explanation. Perhaps one might hope to argue from inference to the best explanation concerning the relations of higher-level to more basic theories in the natural sciences that the entities postulated by psychology should supervene on those of physiology. Or perhaps one might try to draw analogies between non-individualistic theories in psychology and past, unsuccessful theories. These two strategies might meet our methodological strictures on answering the question of whether non-individualistic explanations are viable in a way that an unalloyed appeal to a supervenience principle does not. But philosophical invocations of inference to the best explanation tend to conceal wild leaps supported primarily by ideology. Such considerations must be spelled out into arguments. So far they do not seem very promising.

Take the first strategy. Inductions from the natural sciences to the human sciences are problematic from the start. The problems of the two sorts of sciences look very different, in a multitude of ways. One can, of course, reasonably try to exploit analogies in a pragmatic spirit. But the fact that some given analogy does not hold hardly counts against an otherwise viable mode of explanation. Moreover, there are non-individualistic modes of explanation even in the natural sciences. Geology, physiology, and other parts of biology appeal to entities that are not supervenient on their underlying physical make up. Kind notions in these sciences

[9]Even more generally, I think that epistemic power in philosophy derives largely from reflections on particular implementations of successful cognitive practices. By a cognitive practice, I mean a cognitive enterprise that is stable, that conforms to standard conditions of inter-subjective checkability, and that incorporates a substantial core of agreement among its practitioners. Revisionistic philosophical hypotheses must not, of course, be rejected out of hand. Sometimes, but rarely nowadays, such hypotheses influence cognitive practices by expanding theoretical imagination so as to lead to new discoveries. The changed practice may vindicate the philosophical hypothesis. But the hypothesis waits on such vindication.

(plates, organs, species) presuppose individuative methods that make essential reference to the environment surrounding instances of those kinds.

The second strategy seems even less promising. As it stands, it is afflicted with a bad case of vagueness. Some authors have suggested similarities between vitalism in biology, or action-at-a-distance theories in physics, and non-individualist theories in psychology. The analogies are tenuous. Unlike vitalism, non-individualist psychology does not ipso facto appeal to a new sort of force. Unlike action-at-a-distance theories, it does not appeal to action at a distance. It is true that aspects of the environment that do not differentially affect the physical movement of the protagonists in the thought experiments do differentially affect the explanations and descriptions. This is not, however, because some special causal relation is postulated, but rather because environmental differences affect what kinds of laws obtain, and the way causes and effects are individuated.

Let us now consider a further type of objection to applying the thought experiments to psychology. Since the actual and counterfactual protagonists are so impressively *similar* in so many psychologically relevant ways, can a theoretical language that cuts across these similarities be empirically adequate? The physiological and non-intensional "behavioral" similarities between the protagonists seem to demand similarity of explanation. In its stronger form this objection purports to indicate that non-individualistic mentalistic language has no place in psychology. In its weaker form it attempts to motivate a new theoretical language that attributes intensional content, yet is individualistic. Only the stronger form would establish individualism in psychology. I shall consider it first.

The objection is that the similarities between the protagonists render implausible any theory that treats them differently. This objection is vague or enthymemic. Filling it out tends to lead one back toward the arguments that we have already rejected. On any view, there are several means available (neurophysiology, parts of psychology) for explaining in similar fashion those similarities that are postulated between protagonists in the thought experiments. The argument is not even of the right form to produce a reason for thinking that the differences between the protagonists should not be reflected somewhere in psychological theory—precisely the point at issue.

The objection is often coupled with the remark that non-individualistic explanations would make the parallels between the behavior of the protagonists in the thought experiments "miraculous": explaining the same behavioral phenomena as resulting from different propositional attitudes would be to invoke a "miracle." The rhetoric about miracles can be deflated by noting that the protagonists' "behavior" is not straightforwardly identical, that non-individualistic explanations postulate no special forces, and that there are physical differences in the protagonists' environments that help motivate describing and explaining their activity, at least at one level, in different ways.

The rhetoric about miracles borders on a fundamental misunderstanding of the status of the non-individualistic thought experiments, and of the relation between philosophy and psychology. There is, of course, considerable empirical implausibility, which we might with some exaggeration call "miraculousness," in two person's having identical individualistic physical histories but different thoughts. Most of this implausibility is an artifact of the two-person version of the thought experiments—a feature that is quite inessential. (One may take a single person in two counterfactual circumstances.) This point raises a caution. It is important not to think of the thought experiments as if they were describing actual empirical cases. Let me articulate this remark.

The kinds of a theory, and its principles of individuation, evolve in response to the world as it actually is found to be. Our notions of similarity result from attempts to explain actual cases. They are not necessarily responsive to preconceived philosophical ideals.[10] The kind terms of propositional attitude discourse are responsive to broad, stable similarities in the actual environment that agents are taken to respond to, operate on, and represent. If theory had been frequently confronted with physically similar agents in different environments, it might have evolved different kind terms. But we are so far from being confronted by even rough approximations to global physical similarities between agents that there is little plausibility in imposing individual physical similarity by itself as an ideal sufficient condition for sameness of kind terms throughout psy-

[10]For an interesting elaboration of this theme in an experimental context, see Amos Tversky, "Features of Similarity," *Psychological Review* 84 (1977), pp. 327–352. Cf. also Rosch et al., *op cit.*

chology. Moreover, I think that local physical similarities between the psychologically relevant activities of agents are so frequently intertwined with environmental constancies that a psychological theory that insisted on entirely abstracting from the nature of the environment in choosing its kind terms would be empirically emasculate.

The correct use of counterfactuals in the thought experiments is to explore the scope and limits of the kind notions that have been antecedently developed in attempts to explain actual empirical cases. In counterfactual reasoning we assume an understanding of what our language expresses and explore its application conditions through considering non-actual applications. The counterfactuals in the philosophical thought experiments illumine individuative and theoretical principles to which we are already committed.

The empirical implausibility of the thought experiments is irrelevant to their philosophical point—which concerns possibility, not plausibility. Unlikely but limiting cases are sometimes needed to clarify the modal status of presuppositions that govern more mundane examples. Conversely, the highly counterfactual cases are largely irrelevant to *evaluating* an empirical theory—except in cases (not at issue here) where they present empirical possibilities that a theory counts impossible. To invoke a general philosophical principle, like the supervenience principle, or to insist in the face of the thought experiments that only certain sorts of similarity can be relevant to psychology—without criticizing psychological theory on empirical grounds or showing how the kind notions exhibited by the thought experiments are empirically inadequate—is either to treat counterfactual circumstances as if they were actual, or to fall into *apriorism* about empirical science.

Let us turn to the weaker form of the worry that we have been considering. The worry purports to motivate a new individualistic language of attitude attribution. As I have noted, accepting such a language is consistent with rejecting (pan-) individualism in psychology. There are a variety of levels or kinds of explanation in psychology. Adding another will not alter the issues at stake here. But let us pursue the matter briefly.

There are in psychology levels of individualistic description above the physiological but below the attitudinal that play a role in systematic explanations. Formalistically described computational

processes are appealed to in the attempt to specify an algorithm by which a person's propositional information is processed. I think that the protagonists in our thought experiments might, for some purposes, be said to go through identical algorithms formalistically described. Different information is processed in the "same" ways, at least at this formal level of description. But then might we not want a whole level of description, between the formal algorithm and ordinary propositional attitude ascription, that counts "information" everywhere the same between protagonists in the thought experiments? This is a difficult and complex question, which I shall not attempt to answer here. I do, however, want to mention grounds for caution about supplementing psychology wholesale.

In the first place, the motivation for demanding the relevant additions to psychological theory is empirically weak. In recent philosophical literature, the motivation rests largely on intuitions about Cartesian demons or brains in vats, whose relevance and even coherence have been repeatedly challenged; on preconceptions about the supervenience of the mental on the neural that have no generalized scientific warrant; on misapplications of ordinary observations about causation; and on a sketchy and unclear conception of behavior unsupported by scientific practice.[11] Of course, one may reasonably investigate any hypothesis on no more than an intuitively based hunch. What is questionable is the view that there are currently strong philosophical or scientific grounds for instituting a new type of individualistic explanation.

In the second place, it is easy to underestimate what is involved in creating a relevant individualistic language that would be of genuine use in psychology. Explications of such language have so far been pretty make-shift. It does not suffice to sketch a semantics that says in effect that a sentence comes out true in all worlds that chemically identical protagonists in relevant thought experiments

[11]The most careful and plausible of several papers advocating a new language of individualist explanation is Stephen White, "Partial Character and the Language of Thought," *Pacific Philosophical Quarterly* 63 (1982), pp. 347–365. It seems to me, however, that many of the problems mentioned in the text here and below, beset this advocacy. Moreover, the positive tasks set for the new language are already performed by the actual non–individualist language of psychology. The brain-in-vat intuitions raise very complex issues that I cannot pursue here. I discuss them further in "Cartesian Error and the Objectivity of Perception," *op. cit.*

cannot distinguish. Such an explication gives no clear rules for the *use* of the language, much less a demonstration that it can do distinctive work in psychology. Moreover, explication of the individualistic language (or language component) only for the special case in which the language-user's physiological or (individualistically specified) functional states are held constant, is psychologically useless since no two people are ever actually identical in their physical states.

To fashion an individualist language it will not do to limit its reference to objective properties accessible to perception. For our language for ascribing notions of perceptually accessible physical properties is not individualistic. More generally, as I have argued elsewhere (last *op. cit.* note 1), any attitudes that contain notions for physical objects, events and properties are non-individualistic.[12] The assumptions about objective representation needed to generate the argument are very minimal. I think it questionable whether there is a coherent conception of objective representation that can support an individualistic language of intentional attitude attribution. Advocates of such a language must either explain such a conception in depth, or attribute intentional states that lack objective physical reference.

II.

I have been criticizing arguments for revising the language of psychology to accord with individualism. I have not tried to argue for non-individualistic psychological theories from a standpoint outside of psychology. The heart of my case is the observation that psychological theories, taken literally, are not purely individualistic, that there are no strong reasons for taking them non-literally, and that we currently have no superior standpoint for judging how psychology ought to be done than that of seeing how it *is* done. One can, of course, seek deeper understanding of non-individualistic aspects of psychological theory. Development of such understanding is a multi-faceted task. Here I shall develop

[12]See especially "Intellectual Norms and Foundations of Mind," *op. cit.*, but also "Individualism and the Mental," *op. cit.*, pp. 81–82.

only points that are crucial to my thesis, illustrating them in some detail by reference to one theory.

Ascription of intentional states and events in psychology constitutes a type of individuation and explanation that carries presuppositions about the specific nature of the person's or animal's surrounding environment. Moreover, states and events are individuated so as to set the terms for specific evaluations of them for truth or other types of success. We can judge directly whether conative states are practically successful and cognitive states are veridical. For example, by characterizing a subject as visually representing an X, and specifying whether the visual state appropriately derives from an X in the particular case, we can judge whether the subject's state is veridical. Theories of vision, of belief formation, of memory, learning, decision-making, categorization, and perhaps even reasoning all attribute states that are subject to practical and semantical evaluation *by reference to standards partly set by a wider environment.*

Psychological theories are not themselves evaluative theories. But they often individuate phenomena so as to make evaluation readily accessible *because* they are partly motivated by such judgments. Thus we judge that in certain delimitable contexts people get what they want, know what is the case, and perceive what is there. And we try to frame explanations that account for these successes, and correlative failures, in such a way as to illumine as specifically as possible the mechanisms that underly and make true our evaluations.

I want to illustrate and develop these points by considering at some length a theory of vision. I choose this example primarily because it is a very advanced and impressive theory, and admits to being treated in some depth. Its information-processing approach is congenial with mainstream work in cognitive psychology. Some of its intentional aspects are well understood—and indeed are sometimes conceptually and mathematically far ahead of its formal (or syntactical) and physiological aspects. Thus the theory provides an example of a mentalistic theory with solid achievements to its credit.

The theory of vision maintains a pivotal position in psychology. Since perceptual processes provide the input for many higher cognitive processes, it is reasonable to think that if the theory of vision

treats intentional states non-individualistically, other central parts of cognitive psychology will do likewise. Information processed by more central capacities depends, to a large extent, on visual information.

Certain special aspects of the vision example must be noted at the outset. The arguments that I have previously published against individualism (cf. note 1) have centered on "higher" mental capacities, some of which essentially involve the use of language. This focus was motivated by an interest in the relation between thought and linguistic meaning and in certain sorts of intellectual responsibility. Early human vision makes use of a limited range of representations—representations of shape, texture, depth and other spatial relations, motion, color, and so forth. These representations (percepts) are formed by processes that are relatively immune to correction from other sources of information; and the representations of early vision appear to be fully independent of language. So the thought experiments that I have previously elaborated will not carry over simply to early human vision. (One would expect those thought experiments to be more relevant to social and developmental psychology, to concept learning, and to parts of "higher" cognitive psychology.) But the case against individualism need not center on higher cognitive capacities or on the relation between thought and language. The anti-individualistic conclusions of our previous arguments can be shown to apply to early human vision. The abstract schema which those thought experiments articulate also applies.

The schema rests on three general facts. The first is that what entities in the objective world one intentionally interacts with in the employment of many representational (intentional) types affects the semantical properties of those representational types, what they are, and how we individuate them.[13] A near consequence of this first fact is that there can be slack between, on the one hand, the way a subject's representational types apply to the world, and

[13] 'Representational type' (also 'intentional type') is a relatively theory-neutral term for intentional content, or even intentional state-kinds. Cf. note 3. One could about as well speak of concepts, percepts, and the representational or intentional aspects of thought contents—or of the counterpart states.

on the other, what that person knows about, and how he or she can react to, the way they apply. It is possible for representational types to apply differently, without the person's physical reactions or discriminative powers being different. These facts, together with the fact that many fundamental mental states and events are individuated in terms of the relevant representational types, suffice to generate the conclusion that many paradigmatic mental states and events are not individualistically individuated: they may vary while a person's body and discriminative powers are conceived as constant. For by the second fact one can conceive of the way a person's representational types apply to the objective world as varying, while that person's history, non-intentionally and individualistically specified, is held constant. By the first fact, such variation may vary the individuation of the person's representational types. And by the third, such variation may affect the individuation of the person's mental states and events. I shall illustrate how instances of this schema are supported by Marr's theory of vision.[14]

---

[14]In what follows I make use of the important book *Vision,* by David Marr, (San Francisco: W. H. Freeman and Company, 1982). Marr writes:

> The purpose of these representations is to provide useful descriptions of aspects of the real world. The structure of the real world therefore plays an important role in determining both the nature of the representations that are used and the nature of the processes that derive and maintain them. An important part of the theoretical analysis is to make explicit the physical constraints and assumptions that have been used in the design of the representations and processes . . . (p. 43).
>
> It is of critical importance that the tokens [representational particulars] one obtains [in the theoretical analysis] correspond to real physical changes on the viewed surface; the blobs, lines, edges, groups, and so forth that we shall use must not be artifacts of the imaging process, or else inferences made from their structure backwards to the structures of the surface will be meaningless (p. 44).

Marr's claim that the structure of the real world figures in determining the nature of the representations that are attributed in the theory is tantamount to the chief point about representation or reference that generates our non-individualist thought experiments—the first step in the schema. I shall show that these remarks constitute the central theoretical orientation of the book.

Calling the theory Marr's is convenient but misleading. Very substantial contributions have been made by many others; and the approach has developed rapidly since Marr's death. Cf. for example, Ballard, Hinton, and Sejnowski, "Parallel Vision Computation," *Nature* 306 (November 1983), pp. 21–26. What I say about Marr's book applies equally to more recent developments.

Marr's theory subsumes three explanatory enterprises: (a) a theory of the computation of the information, (b) an account of the representations used and of the algorithms by which they are manipulated, and (c) a theory of the underlying physiology. Our primary interest is in the first level, and in that part of the second that deals with the individuation of representations. Both of these parts of the theory are fundamentally intentional.

The theory of the computation of information encompasses an account of what information is extracted from what antecedent resources, and an account of the reference-preserving "logic" of the extraction. These accounts proceed against a set of biological background assumptions. It is assumed that visual systems have evolved to solve certain problems forced on them by the environment. Different species are set different problems and solve them differently. The theory of human vision specifies a general information processing problem—that of generating reliable representations of certain objective, distal properties of the surrounding world on the basis of proximal stimulations.

The human visual system computes complex representations of certain visible properties, on the basis of light intensity values on retinal images. The primary visible properties that Marr's theory treats are the shapes and locations of things in the world. But various other properties—motion, texture, color, lightness, shading—are also dealt with in some detail. The overall computation is broken down into stages of increasing complexity, each containing modules that solve various subproblems.

The theory of computation of information clearly treats the visual system as going through a series of intentional or representational states. At an early stage, the visual system is counted as representing objective features of the physical world.[15] There is no

[15]It is an interesting question when to count the visual system as having gone intentional. I take it that information is in a broad sense, carried by the intensity values in the retinal image; but I think that this is too early to count the system as intentional or symbolic. I'm inclined to agree with Marr that where zero-crossings from different sized filters are checked against one another (cf. Example 1), it is reasonable to count visual processes as representational of an external physical reality. Doing so, however, depends on seeing this stage as part of the larger system in which objective properties are often discriminated from subjective artifacts of the visual system.

other way to treat the visual system as solving the problem that the theory sees it as solving than by attributing intentional states that represent objective, physical properties.

More than half of Marr's book is concerned with developing the theory of the computation of information and with individuating representational primitives. These parts of the theory are more deeply developed, both conceptually and mathematically, than the account of the algorithms. This point is worth emphasizing because it serves to correct the impression, often conveyed in recent philosophy of psychology, that intentional theories are regressive and all of the development of genuine theory in psychology has been proceeding at the level of purely formal, "syntactical" transformations (algorithms) that are used in cognitive systems.

I now want, by a series of examples, to give a fairly concrete sense of how the theory treats the relation between the visual system and the physical environment. Understanding this relation will form essential background for understanding the non-individualistic character of the theory. The reader may skip the detail and still follow the philosophical argument. But the detail is there to support the argument and to render the conception of explanation that the argument yields both concrete and vivid.

Initially, I will illustrate two broad points. The *first* is that the theory makes essential reference to the subject's distal stimuli and makes essential assumptions about contingent facts regarding the subject's physical environment. Not only do the basic questions of the theory refer to what one sees under normal conditions, but the computational theory and its theorems are derived from numerous explicit assumptions about the physical world.

The *second* point to be illustrated is that the theory is set up to explain the reliability of a great variety of processes and sub-processes for acquiring information, at least to the extent that they are reliable. Reliability is presupposed in the formulations of the theory's basic questions. It is also explained through a detailed account of how in certain specified, standard conditions, veridical information is derived from limited means. The theory explains not merely the reliability of the system as a whole, but the reliability of various stages in the visual process. It begins by assuming that we see certain objective properties and proceeds to explain particular successes by framing conditions under which success would be ex-

pected (where the conditions are in fact typical). Failures are explained primarily by reference to a failure of these conditions to obtain. To use a phrase of Bernie Kobes, the theory is not success-neutral. The explanations and, as we shall later see, the kinds of theory presuppose that perception and numerous subroutines of perception are veridical in normal circumstances.

Example 1: In an early stage of the construction of visual representation, the outputs of channels or filters that are sensitive to spatial distributions of light intensities are combined to produce representations of local contours, edges, shadows, and so forth. The filters fall into groups of different sizes, in the sense that different groups are sensitive to different bands of spatial frequencies. The channels are primarily sensitive to sudden intensity changes, called "zero-crossings," at their scales (within their frequency bands). The theoretical question arises: How do we combine the results of the different sized channels to construct representations with physical meaning—representations that indicate edge segments or local contours in the external physical world? There is no *a priori* reason why zero-crossings obtained from different sized filters should be related to some one physical phenomenon in the environment. There is, however, a physical basis for their being thus related. This basis is identified by *the constraint of spatial localization.* Things in the world that give rise to intensity changes in the image, such as changes of illumination (caused by shadows, light sources) or changes in surface reflectance (caused by contours, creases, and surface boundaries), are spatially localized, not scattered and not made up of waves. Because of this fact, if a zero-crossing is present in a channel centered on a given frequency band, there should be a corresponding zero-crossing at the same spatial location in larger-scaled channels. If this ceases to be so at larger scales, it is because a) two or more local intensity changes are being averaged together in the larger channel (for example, the edges of a thin bar may register radical frequency changes in small channels, but go undetected in larger ones); or b) because two independent physical phenomena are producing intensity changes in the same area but at different scales (for example, a shadow superimposed on a sudden reflectance change; if the shadow is located in a certain way, the positions of the zero-crossings may not make possible a separation of the two physical phenomena). Some of these exceptions are sufficiently rare that the visual system need not and does not account for them—thus allowing for possible illusions; others are reflected in complications of the basic assumption that follows. The spatial coincidence constraint yields *the spatial coincidence assumption:*

If a zero-crossing segment is present in a set of independent channels over a contiguous range of sizes, and the segment has the same position and orienta-

tion in each channel, then the set of such zero-crossing segments indicates the presence of an intensity change in the image that is due to a single physical phenomenon (a change in reflectance, illumination, depth, or surface orientation).

Thus the theory starts with the observation that physical edges produce roughly coincident zero-crossings in channels of neighboring sizes. The spatial coincidence assumption asserts that the coincidence of zero-crossings of neighboring sizes is normally sufficient evidence of a real physical edge. Under such circumstances, according to the theory, a representation of an edge is formed.[16]

Example 2: Because of the laws of light and the way our eyes are made, positioned, and controlled, our brains typically receive similar image signals originating from two points that are fairly similarly located in the respective eyes or images, at the same horizontal level. If two objects are separated in depth from the viewer, the relative positions of their image signals will differ in the two eyes. The visual system determines the distance of physical surfaces by measuring the angular discrepancy in position (disparity) of the image of an object in the two eyes. This process is called stereopsis. To solve the problem of determining distance, the visual system must select a location on a surface as represented by one image, identify the same location in the other image, and measure the disparity between the corresponding image points. There is, of course, no *a priori* means of matching points from the two images. The theory indicates how correct matches are produced by appealing to three *Physical Constraints* (actually the first is not made explicit, but is relied upon): (1) the two eyes produce similar representations of the same external items; (2) a given point on a physical surface has a unique position in space at any given time; (3) matter is cohesive—separated into objects, the surfaces of which are usually smooth in the sense that surface variation is small compared to overall distance from the observer. These three physical constraints are rewritten as three corresponding *Constraints on Matching:* (1) two representational elements can match if and only if they normally could have arisen from the same physical item (for example, in stereograms, dots match dots rather than bars); (2) nearly always, each representational element can match only one element from the other image (exceptions occur when two markings lie along the line of sight of one eye but are separately visible by the other—causing illusions); (3) disparity varies smoothly almost everywhere (this derives from physical constraint (3) because that constraint implies that the distance

[16]Marr, *op. cit.*, pp. 68–70; cf. also Marr and Hildreth, "Theory of Edge Detection," *Proceedings of Royal Society of London* B 207 (1980), pp. 187–217, where the account is substantially more detailed.

to the visible surface varies, approximately continuously except at object boundaries, which occupy a small fraction of the area of an image). Given suitable precisifications, these matching constraints can be used to prove the *Fundamental Theorem of Stereopsis:*

> If a correspondence is established between physically meaningful representational primitives extracted from the left and right images of a scene that contains a sufficient amount of detail (roughly 2% density for dot stereograms), and if the correspondence satisfies the three matching constraints, then that correspondence is physically correct—hence unique.

> The method is again to identify general physical conditions that give rise to a visual process, then to use those conditions to motivate constraints on the form of the process that, when satisfied, will allow the process to be interpreted as providing reliable representations of the physical environment.[17]

These examples illustrate theories of the computation of information. The critical move is the formulation of general physical facts that limit the interpretation of a visual problem enough to allow one to interpret the machinations of the visual system as providing a unique and veridical solution, at least in typical cases. The primary aim of referring to contingent physical facts and properties is to enable the theory to explain the visual system's reliable acquisition of information about the physical world: to explain the success or veridicality of various types of visual representation. So much for the first two points that we set out to illustrate.

I now turn to a *third* that is a natural corollary of the second, and that will be critical for our argument that the theory is non-individualistic: the information carried by representations—their intentional content—is individuated in terms of the specific distal causal antecedents in the physical world that the information is about and that the representations normally apply to. The individuation of the intentional features of numerous representations depends on a variety of physical constraints that our knowledge of the external world gives us. Thus the individuation of intentional

[17]Marr, *op. cit.*, pp. 111–116; Marr and Poggio, "A Computational Theory of Human Stereo Vision," *Proceedings of Royal Society of London* B 204 (1979), pp. 301–328. Marr, *op. cit.*, pp. 205–212; Shimon Ullman, *The Interpretation of Visual Motion,* (Cambridge, Mass.: MIT Press, 1979).

content of representational types, presupposes the verdicality of perception. Not only the explanations, but the intentional kinds of the theory presuppose contingent facts about the subject's physical environment.

Example 3: In building up informational or representational primitives in the primal sketch, Marr states six general physical assumptions that constrain the choice of primitives. I shall state some of these to give a sense of their character: (a) the visible world is composed of smooth surfaces having reflectance functions whose spatial structure may be complex; (b) markings generated on a surface by a single process are often arranged in continuous spatial structures—curves, lines, etc.; (c) if direction of motion is discontinuous at more than one point—for example, along a line—then an object boundary is present. These assumptions are used to identify the physical significance of—the objective information normally given by—certain types of patterns in the image. The computational theory states conditions under which these primitives form to carry information about items in the physical world (Marr, *op. cit.*, pp. 44–71). The theory in Example 1 is a case in point: conditions are laid down under which certain patterns may be taken as representing an objective physical condition; as being edge, boundary, bar, or blob detectors. Similar points apply for more advanced primitives.

Example 4: In answering the question "what assumptions do we reasonably and actually employ when we interpret silhouettes as three-dimensional shapes?" Marr motivates a central representational primitive by stating physical constraints that lead to the proof of a theorem. *Physical Constraints:* (1) Each line of sight from the viewer to the object grazes the object's surface at exactly one point. (2) Nearby points on the contour in an image arise from nearby points on the contour generator on the viewed object. (That is, points that appear close together in the image actually are close together on the object's surface.) (3) The contour generator lies wholly in a single plane. Obviously, these are conditions of perception that may fail, but they are conditions under which we seem to do best at solving the problem of deriving three-dimensional shape descriptions from representations of silhouettes. *Definition:* A *generalized cone* is a three-dimensional object generated by moving a cross section along an axis; the cross section may vary smoothly in size, but its shape remains the same. (For example footballs, pyramids, legs, stalagmites are or approximate generalized cones.) *Theorem:* If the surface is smooth and if physical constraints (1)–(3) hold for all distant viewing positions in any one plane, then the viewed surface is a generalized cone. The theorem indicates a natural connection between generalized cones and the

> imaging process. Marr infers from this, and from certain psychophysical evidence, that representations of generalized cones—that is, representations with intentional content concerning, generalized cones—are likely to be fundamental among our visual representations of three-dimensional objects (Marr, *op. cit.*, pp. 215–225).

Throughout the theory, representational primitives are selected and individuated by considering specific, contingent facts about the physical world that typically hold when we succeed in obtaining veridical visual information about that world. The information or content of the visual representations is always individuated by reference to the physical objects, properties, or relations that are seen. In view of the success-orientation of the theory, this mode of individuation is grounded in its basic methods. If theory were confronted with a species of organism reliably and successfully interacting with a different set of objective visible properties, the representational types that the theory would attribute to the organism would be different, regardless of whether an individual organism's physical mechanisms were different.

We are now in a position to argue that the theory is not individualistic: (1) The theory is intentional. (2) The intentional primitives of the theory and the information they carry are individuated by reference to contingently existing physical items or conditions by which they are normally caused and to which they normally apply. (3) So if these physical conditions and, possibly, attendant physical laws were regularly different, the information conveyed to the subject and the intentional content of his or her visual representations would be different. (4) It is not incoherent to conceive of relevantly different physical conditions and perhaps relevantly different (say, optical) laws regularly causing the same non-intentionally, individualistically individuated physical regularities in the subject's eyes and nervous system. It is enough if the differences are small; they need not be wholesale. (5) In such a case (by (3)) the individual's visual representations would carry different information and have different representational content, though the person's whole non-intentional physical history (at least up to a certain time) might remain the same. (6) Assuming that some perceptual states are identified in the theory in terms of their informational or intentional content, it follows that individualism is not true for the theory of vision.

I shall defend the argument stepwise. I take it that the claim that the theory is intentional is sufficiently evident. The top levels of the theory are explicitly formulated in intentional terms. And their method of explanation is to show how the problem of arriving at certain veridical representations is solved.

The second step of the argument was substantiated through Examples 3 and 4. The intentional content of representations of edges or generalized cones is individuated in terms of *specific* reference to those very contingently instantiated physical properties, on the assumption that those properties normally give rise to veridical representations of them.

The third step in our argument is supported both by the way the theory individuates intentional content (cf. the previous paragraph and Examples 3 and 4), and by the explanatory method of the theory (cf. the second point illustrated above, and Examples 1–2). The methods of individuation and explanation are governed by the assumption that the subject has adapted to his or her environment sufficiently to obtain veridical information from it under certain normal conditions. If the properties and relations that *normally* caused visual impressions were regularly different from what they are, the individual would obtain different information and have visual experiences with different intentional content. If the regular, law-like relations between perception and the environment were different, the visual system would be solving different information-processing problems; it would pass through different informational or intentional states; and the explanation of vision would be different. To reject this third step of our argument would be to reject the theory's basic methods and questions. But these methods and questions have already borne fruit, and there are presently no good reasons for rejecting them.

I take it that step four is a relatively unproblematic counterfactual. There is no metaphysically necessary relation between individualistically individuated processes in a person's body and the causal antecedents of those processes in the surrounding world.[18] (To reject this step would be self-defeating for the individualist.) If

[18]As I have intimated above, I doubt that all biological, including physiological, processes and states in a person's body are individualistically individuated. The failures of individualism for these sciences involve different, but related considerations.

the environmental conditions were different, the same proximal *visual* stimulations could have regularly had different distal causes. In principle, we can conceive of some regular variation in the distal causes of perceptual impressions with no variation in a person's individualistically specified physical processes, even while conceiving the person as *well adapted* to the relevant environment—though, of course, not uniquely adapted.

Steps three and four, together with the unproblematic claim that the theory individuates some perceptual states in terms of their intentional content or representational types, entail that the theory is non-individualistic.

Steps two and three are incompatible with certain philosophical approaches that have no basis in psychological theory. One might claim that the information content of a visual representation would remain constant even if the physical conditions that lead to the representation were regularly different. It is common to motivate this claim by pointing out that one's visual representations remain the same, whether one is perceiving a black blob on a white surface or having an eidetic hallucination of such a blob. So, runs the reasoning, why should changing the distal causes of a perceptual representation affect its content? On this view, the content of a given perceptual representation is commonly given as that of "the distal cause of *this* representation," or "the property in the world that has *this* sort of visual appearance." The content of these descriptions is intended to remain constant between possible situations in which the micro-physical events of a person's visual processes remain the same while distal causes of those processes are regularly and significantly different. For it is thought that the representations themselves (and our experiences of *them*) remain constant under these circumstances. So as the distal antecedents of one's perceptual representations vary, the reference of those representations will vary, but their intentional content will not.[19]

[19]Descartes went further in the same direction. He thought that the perceptual system, and indeed the intellect, could not make a mistake. Mistakes derived from the will. The underlying view is that we primarily perceive or make perceptual reference to our own perceptions. This position fails to account plausibly for various visual illusions and errors that precede any activity of the will, or even intellect. And the idea that perceptions are in general what we make perceptual reference to has little to

There is more wrong with this line than I have room to develop here. I will mention some of the more straightforward difficulties. In the first place, the motivation from perceptual illusion falls far short. One is indeed in the same perceptual state whether one is seeing or hallucinating. But that is because the intentional content of one's visual state (or representation) is individuated against a background in which the relevant state is *normally* veridical. Thus the fact that one's percepts or perceptual states remain constant between normal perception and hallucinations does not even tend to show that the intentional visual state remains constant between circumstances in which different physical conditions are the normal antecedents of one's perceptions.

Let us consider the proposals for interpreting the content of our visual representations. In the first place both descriptions ('the distal cause of *this* representation' *et al.*) are insufficiently specific. There are lots of distal causes and lots of things that might be said to appear "thus" (for example, the array of light striking the retina as well as the physical surface). We identify the relevant distal cause (and the thing that normally appears thus and so) as the thing that we actually see. To accurately pick out the "correct" object with one of these descriptions would at the very least require a more complex specification. But filling out the descriptive content runs into one or both of two difficulties: either it includes kinds that are tied to a specific environment ('the convex, rough textured object that is causing this representation'). In such case, the description is still subject to our argument. For these kinds are individuated by reference to the empirical environment. Or it complicates the constraints on the causal chain to the extent that the complications cannot plausibly be attributed to the content of processes in the early visual system.

Even in their unrevised forms, the descriptions are over-intellectualized philosophers' conceits. It is extremely implausible and em-

---

recommend it and, nowadays, little influence. The natural and, I think, plausible view is that we have visual representations that specify external properties specifically, that these representations are pre-doxastic in the sense they are not themselves objects of belief, and that they sometimes fail to represent correctly what is before the person's eyes: when they result from abnormal processes.

pirically without warrant to think that packed into every perceptual representation is a distinction between distal cause and experiential effect, or between objective reality and perceptual appearance. These are distinctions developed by reflecting on the ups and downs of visual perception. They do not come in at the ground, animal level of early vision.

A further mistake is the view that our perceptual representations never purport to specify particular physical properties *as such,* but only via some relation they bear to inner occurrences, which are directly referred to. (Even the phrase 'the convex object causing this percept' invokes a specification of objective convexity as such.) The view will not serve the needs of psychological explanation as actually practiced. For the descriptions of information are too inspecific to account for specific successes in solving problems in retrieving information about the actual, objective world.

The best empirical theory that we have individuates the intentional content of visual representations by specific reference to specific physical characteristics of visible properties and relations. The theory does not utilize complicated, self-referential, attributively used role descriptions of those properties. It does not individuate content primarily by reference to phenomenological qualities. Nor does it use the notions of cause or appearance in specifying the intentional content of early visual representations.[20]

The second and third steps of our argument are incompatible

[20]Of course, at least in the earliest stages of visual representation, there are analogies between qualitative features of representations in the experienced image and the features that those representations represent. Representations that represent bar segments are bar-shaped, or have some phenomenological property that strongly tempts us to call them "bar-shaped." Similarly for blobs, dots, lines and so forth. (Marr and Hildreth, *op. cit.,* p. 211, remark on this dual aspect of representations.) These "analogies" are hardly fortuitous. Eventually they will probably receive rigorous psychophysical explanations. But they should not tempt one into the idea that visual representations in general make reference to themselves, much less into the idea that the content of objective representation is independent of empirical relations between the representations and the objective entities that give rise to them. Perhaps these qualitative features are constant across all cases where one's bodily processes, non-intentionally specified, are held constant. But the information they carry, their intentional content, may vary with their causal antecedents and causal laws in the environment.

with the claim that the intentional content of visual representations is determined by their "functional role" in each person's system of dispositions, non-intentionally and individualistically specified. This claim lacks any warrant in the practice of the science. In the first place, the theory suggests no reduction of the intentional to the non-intentional. In the second, although what a person can do, non-visually, constitutes evidence for what he or she can see, there is little ground for thinking that either science or common sense takes an individual person's non-visual abilities fully to determine the content of his or her early visual experience. A person's dispositions and beliefs develop by adapting to what the person sees. As the person develops, the visual system (at least at its more advanced stages—those involving recognition) and the belief and language systems affect each other. But early vision seems relatively independent of these non-visual systems. A large part of learning is accommodating one's dispositions to the information carried by visual representations. Where there are failures of adaptation, the person does not know what the visual apparatus is presenting to him or her. Yet the presentations are there to be understood.

III

There is a general argument that seems to me to show that a person's non-intentional dispositions could not fix (individuate) the intentional content of the person's visual presentations. The argument begins with a conception of objectivity. As long as the person's visual presentations are of public, objective objects, properties, or relations, it is possible for the person to have mistaken presentations. Such mistakes usually arise for a single sensory modality—so that when dispositions associated with other modalities (for example, touch) are brought into play, the mistake is rectified. But as long as the represented object or property is objective and physical, it is in principle possible, however unlikely, that there be a confluence of illusions such that all an individual person's sensory modalities would be fooled and all of the person's non-intentional dispositions would fail to distinguish between the normal condition and the one producing the mistaken sensory representations. This is our first assumption. In the argument, we shall employ a corollary: our concept of objectivity is such that no

one objective entity that we visually represent is such that it must vary with, or be typed so as necessarily to match exactly, an individual's proximal stimuli and discriminative abilities. The point follows from a realistic, and even from a non-subjectivistic, view of the objects of sight.[21]

We argued earlier that intentional representational types are not in general individuated purely in terms of an attributive role-description of a causal relation, or a relation of appearance-similarity, between external objects and qualitative perceptual representatives of them. For present purposes, this is our second assumption: some objective physical objects and properties are visually represented as such; they are specifically specified.

Third, in order to be empirically informative, some visual representations that represent objective entities as such must have the representational characteristics that they have partly *because* instances regularly enter into certain relations with those objective entities.[22] Their carrying information, their having objective intentional content, consists partly in their being the normal causal products of objective entities. And their specific intentional content depends partly on their being the normal products of the specific

---

[21]There is no need to assume that the abnormal condition is unverifiable. Another person with relevant background information might be able to infer that the abnormal condition is producing a perceptual illusion. In fact, another person with different dispositions might even be able to perceive the difference.

[22]Not all perceptual representations that specify objective entities need have their representational characteristics determined in this way. The representational characters of *some* visual representations (or states) may depend on the subject's background theory or primarily on interaction among other representations. There are hallucinations of purple dragons. (Incidentally, few if any of the perceptual representations—even the conscious perceptual representations—discussed in Marr's theory depend in this way on the subject's conceptual background.) Here, I assume only that *some* visual representations acquire their representational characters through interaction. This amounts to the weak assumption that the formation of some perceptual representations is *empirical.*

Some of the interaction that leads to the formation and representational characters of certain innate perceptual tendencies (or perhaps even representations) may occur in the making of the species, not in the learning histories of individuals. Clearly this complication could be incorporated into a generalization of this third premise—without affecting the anti-individualistic thrust of the argument.

objective entities that give rise to them. That is why we individuate intentional visual representations in terms of the objective entities that they normally apply to, for members of a given species. This is the core of truth in the slogan, sometimes misapplied I think, that mistakes presuppose a background of veridicality.

The assumptions in the three preceding paragraphs enable us to state a general argument against individualism regarding visual states. Consider a person P who normally correctly perceives instances of a particular objective visible property *O*. In such cases, let the intentional type of P's perceptual representation (or perceptual state) be *O'*. Such perceptual representations are normally the product of interaction with instances of *O*. But imagine that for P, perceptual representations typed *O'* are on some few occasions the product of instances of a different objective property *C*. On such occasions, P mistakenly sees an instance of *C* as an *O;* P's perceptual state is of type *O'*. We are assuming that *O'* represents any instance of *O* as such (as an *O*), in the sense of our second premise, not merely in terms of some attributive role description. Since *O'* represents an objective property, we may, by our first premise, conceive of P as lacking at his or her disposal (at every moment up to a given time) any means of discriminating the instances of *C* from instances of *O*.

Now hold fixed both P's physical states (up to the given time) and his or her discriminative abilities, non-intentionally and individualistically specified. But conceive of the world as lacking *O* altogether. Suppose that the optical laws in the counterfactual environment are such that the impressions on P's eyes and the normal causal processes that lead to P's visual representations are explained in terms of *C*'s (or at any rate, in terms of some objective; visible entities other than instances of *O*). Then by our third premise, P's visual representation (or visual state) would not be of intentional type *O'*. At the time when in the actual situation P is misrepresenting a *C* as an *O*, P may counterfactually be perceiving something (say, a *C*) correctly (as a *C*)—if the processes that lead to that visual impression are normal and of a type that normally produces the visual impression that P has on that occasion. So the person's intentional visual states could vary while his or her physical states and non-intentionally specified discriminative abilities remained constant.

The first premise and the methodology of intentional-content individuation articulated in the third premise entail the existence of examples. Since examples usually involve shifts in optical laws, they are hard to fill out in great detail. But it is easiest to imagine concrete cases taken from early but still conscious vision. These limit the number of an individual's dispositions that might be reasonably thought to bear on the content of his or her visual states. Early vision is relatively independent of linguistic or other cognitive abilities. It appears to be relatively modular.

Suppose that the relevant visible entities are very small and not such as to bear heavily on adaptive success. An *O* may be a shadow of a certain small size and shape on a gently contoured surface. A *C* may be a similarly sized, shallow crack. In the actual situation P sees *O*'s regularly and correctly as *O*'s: P's visual representations are properly explained and specified as shadow representations of the relevant sort. We assume that P's visual and other discriminative abilities are fairly normal. P encounters *C*'s very rarely and on those few occasions not only misperceives them as *O*'s, but has no dispositions that would enable him or her to discriminate those instances from *O*'s. We may assume that given P's actual abilities and the actual laws of optics, P would be capable, in ideal circumstances, of visually discriminating some instances of *C*'s (relevantly similar cracks) from instances of *O* (the relevant sort of shadows). But our supposition is that in the actual cases where P is confronted by instances of *C*'s, the circumstances are not ideal. All P's abilities would not succeed in discriminating those instances of relevant cracks, in those circumstances, from instances of relevant shadows. P may not rely on touch in cases of such small objects; or touch may also be fooled. P's ability to have such mistaken visual states is argued for by the objectivity premise.

In the counterfactual case, the environment is different. There are no instances of the relevant shadows visible to P; and the laws of optics differ in such a way that P's physical visual stimulations (and the rest of P's physical makeup) are unaffected. Suppose that the physical visual stimulations that in the actual case are derived from instances of *O*—at the relevant sort of shadows—are counterfactually caused by and explained in terms of *C*'s, relevantly sized cracks. Counterfactually, the cracks take the places of the shadows. On the few occasions where, in the actual case, P misperceives

shadows as cracks, P is counterfactually confronted with cracks; and the optical circumstances that lead to the visual impressions on those occasions are, we may suppose, normal for the counterfactual environment.[23] On such counterfactual occasions, P would be visually representing small cracks as small cracks. P would never have visual representations of the relevant sort of shadows. One can suppose that even if there were the relevant sort of shadows in the counterfactual environment, the different laws of optics in that environment would not enable P ever to see them. But since P's visual states would be the normal products of normal processes and would provide as good an empirical basis for learning about the counterfactual environment as P has for learning about the actual environment, it would be absurd to hold that (counterfactually) P misperceives the prevalent cracks as shadows on gently contoured surfaces. Counterfactually, P correctly sees the cracks as cracks. So P's intentional perceptual states differ between actual and counterfactual situations. This general argument is independent of the theory of vision that we have been discussing. It supports and is further supported by that theory.

IV.

Although the theory of vision is in various ways special, I see no reason why its non-individualistic methods will not find analogs in other parts of psychology. In fact, as we noted, since vision provides intentional input for other cognitive capacities, there is reason to think that the methods of the theory of vision are presupposed by other parts of psychology. These non-individualistic

[23]What of the non-intentionally specified dispositions that in the actual environment (given the actual laws of optics) would have enabled P to discriminate *C*'s from *O*'s in ideal circumstances? In the counterfactual environment, in view of the very different optical laws and different objects that confront P, one can suppose that these dispositions have almost any visual meaning that one likes. These dispositions would serve to discriminate *C*'s from some other sort of entity. In view of the objectivity premise, the non-intentional dispositions can always be correlated with different, normal antecedent laws and conditions—in terms of which their intentional content may be explained.

The argument of this section is developed in parallel but different ways in "Cartesian Error and the Objectivity of Perception," *op. cit.*

methods are grounded in two natural assumptions. One is that there are psychological states that represent, or are about, an objective world. The other is that there is a scientific account to be given that presupposes certain successes in our interaction with the world (vision, hearing, memory, decision, reasoning, empirical belief formation, communication, and so forth), and that explains specific successes and failures by reference to these states.

The two assumptions are, of course, interrelated. Although an intention to eat meat is "conceptually" related to eating meat, the relation is not one of entailment in either direction, since the representation is about an objective matter. An individual may be, and often is, ignorant, deluded, misdirected, or impotent. The very thing that makes the non-individualistic thought experiments possible—the possibility of certain sorts of ignorance, failure, and misunderstanding—helps make it possible for explanations using non-individualistic language to be empirically informative. On the other hand, as I have argued above, some successful interaction with an objective world seems to be a precondition for the objectivity of some of our intentional representations.

Any attempt to produce detailed accounts of the relations between our attitudes and the surrounding world will confront a compendium of empirically interesting problems. Some of the most normal and mundane successes in our cognitive and conative relations to the world must be explained in terms of surprisingly complicated intervening processes, many of which are themselves partly described in terms of intentional states. Our failures may be explained by reference to *specific* abnormalities in operations or surrounding conditions. Accounting for environmentally specific successes (and failures) is one of the tasks that psychology has traditionally set itself.

An illuminating philosophy of psychology must do justice not only to the mechanistic elements in the science. It must also relate these to psychology's attempt to account for tasks that we succeed and fail at, *where these tasks are set by the environment and represented by the subject him- or herself*. The most salient and important of these tasks are those that arise through relations to the natural and social worlds. A theory that insists on describing the states of human beings *purely* in terms that abstract from their relations to any specific environment cannot hope to provide a completely satisfying

explanation of our accomplishments. At present our best theories in many domains of psychology do not attempt such an abstraction. No sound reason has been given for thinking that the non-individualistic language that psychology now employs is not an appropriate language for explaining these matters, or that explanation of this sort is impossible.

TYLER BURGE

University of California, Los Angeles