

24 Psychology and the Environment: Reply to Chomsky

Noam Chomsky's contributions to the study of language have included several important contributions to philosophy. Most of these contributions—notably ones that are firmly based in his work in linguistics—have been objected to by prominent philosophers. I believe that most of Chomsky's philosophical positions have aged better than the objections. In particular, his opposition to behaviorism and his defense of the methods of studying syntax as a part of psychology turned philosophy, psychology, and linguistics in more fruitful directions. His accounts of the relevance of syntax and phonology to unconscious psychological states, of the large role of innateness in language development, of the universality of many specific linguistic structures, and of the modular character of linguistic abilities—all these have helped provide a specific shape to our understanding of mind. Chomsky's contribution to this volume is presented in an exploratory spirit that leaves open (see note 36) whether and where we disagree on a number of issues. I would like to make progress on clarifying the extent and nature of disagreement, in hopes that this will lead to progress on the issues themselves. Chomsky's essay covers a lot of ground. To sharpen my focus I would like to begin by indicating a number of further points on which we agree.

I

We agree that eliminativism about mental/psychological kinds and structures is not a serious possibility even for science, much less common sense.¹ The progress of cognitive psychology and modern linguistics—among other things—indicates that this idea has gotten more play than it deserves.

We agree that the primary topic of our discussions is meaning and belief, not ordinary beliefs about meaning and belief. Intuitive beliefs about meaning and belief are relevant to understanding these topics. But they are fallible. There is room for discovery of unfamiliar kinds and structures that enter into the explanation or illumination of meaning and belief. Science must make its own way, even where it begins by testing and developing commonsense assumptions.

We agree that many actual linguistic and other psychological structures are inaccessible to conscious introspection, intuitive judgment, and philosophical reflection. The details of psychology—especially the unconscious structures and the psychological and computational mechanisms—are largely beyond the reach of immediate intuitive judgments. They must be discovered through theorizing that appeals to a variety of types of evidence.

There is a difference in emphasis between us here. I believe that whereas nearly all of the structures of syntax are unfamiliar to untutored preconceptions, a larger—but still

minority—proportion of psychological kinds that are relevant to understanding belief and meaning are familiar. This is because a relatively larger proportion of the structure and psychology of meaning and belief is accessible to conscious applicational and introspective understanding. Part of being a critically rational being is having some immediate cognitive control over one's beliefs and meaning. Commonly such control extends to an ability to understand the meanings and beliefs of others. How much of these subject matters is inaccessible is open to investigation. I think, however, that there is no reason whatever to believe that our basic framework for describing and explaining human belief and meaning will be substantially altered or overturned, as opposed to supplemented and enriched.

We agree that there is a large and psychologically specific innate component in our linguistic and other psychological abilities. Many of the basic structures of reasoning, perception, language, and social imitation are probably innate and universal to the species.

We agree that belief-desire-intention psychological explanation of most human action is unsystematic, highly contextual, and at present more the province of common sense than systematic science. We may differ about the cognitive value of common sense, or about how substantial a contribution ordinary belief-desire-intention psychology will make to scientific psychology. I think that human beings have quite a lot of genuine knowledge about the mind through ordinary, nonsystematic judgments. And I believe that some recognizable though vast refinement and elaboration of the belief-desire-intention model will find a place in systematic psychology. But we are in accord about where scientific progress has been made. Most scientific progress in psychology has centered on explaining abilities involved in modular or relatively simple psychological subsystems. We are much further along in explaining the structure of various particular competences than in providing detailed, systematic, law-based explanation of performance—particularly intentional action—except in narrowly confined circumstances.

We agree that there is no decisive philosophical objection to the systematic study of meaning on grounds that “meaning is holistic.” We may differ in our reasons. Chomsky concedes that the slogan applies to the sciences, but denies that it is relevant to natural language. I think that empirical confirmation, both in science and in common sense, is with some qualifications relatively holistic, but that meaning is not to be understood purely or fundamentally in terms of confirmation. I believe that the work on language-world relations (“reference”) and the relation between such relations and meaning count strongly against the positivist account of meaning in terms of confirmation. Perhaps ordinary methodology in lexical semantics does as well. Holism has been enormously overplayed in philosophical accounts of meaning and reference—in accounts both of science and of common sense.

II

In exploring areas of difference, I would like to begin by making some remarks about the terms ‘internalism’ and ‘individualism’. I will sometimes use the former term in this response, though I prefer the latter. It is hard to tell how much disagreement there is here. There is no explicit statement of what internalism is in Chomsky’s essay. Some of his remarks or arguments for “internalism” do not directly connect with my objections to the view I designate with the term, or with the term “individualism.” So some apparent disagreement may not be real.

Chomsky’s I-languages are idiolects—languages of individual speakers. Anti-individualism, even social anti-individualism, does not presuppose the existence of public languages and is perfectly free to focus on individuals’ idiolects.² I think that there is much more to the idea of public, shared linguistic meanings and norms than Chomsky does. But I do not think that anything fundamental to my views hinges on this point. It may well be that the study of syntactic and even semantic structures can, for now, reasonably concentrate on the study of I-languages—though I think that how speakers allow their idiolects to depend on one another is a matter of deep importance. I see no reason why such interdependence cannot be studied systematically and successfully. I think that there is no question that reference and even meaning in individual languages often depend on reliance on others.

Sometimes it appears that Chomsky means by an ‘internalist’ theory one that studies the internal states of individuals. On this understanding of the term ‘internalist’, an externalist theory would include only the study of relations between individuals and the individuals’ environment, including perhaps other individuals of the same kind.³

I think that there are worthwhile systematic studies of relations between individuals and their environment—studies of semantic reference, of social relation, and of other matters. I shall return to this point. But anti-individualism, as I understand it, does not depend on these claims. Internalism or individualism, as I understand it, is not simply a claim that psychology studies the internal states of individuals. Even if I were to agree for the sake of argument that *all* psychological and linguistic theory focused on the internal states of individuals—thus *accepting* internalism in the sense I have found in the passages from Chomsky cited above—I would hold that internalism, *as I understand it*, is mistaken.

Anti-individualism or externalism, in my sense, need not affect the way psychology studies the structures and mechanisms of internal psychological states. I am happy to agree that all psychological states, properly so-called, are “in” the individual’s mind. Anti-individualism is about the nature of “internal” psychological states.

Anti-individualism is the view that an individual’s being in a significant range of particular intentional psychological states (beliefs, understandings, and so on) necessitates

that the individual bear certain causal, functional, or historical relations to an environment beyond the individual. Further, the natures and identities of those states are constitutively dependent on certain relations between individual and environment. The intentional psychological states that require the individual's being in such relations may themselves be internal to the individual—genuine states of the individual. These states are not themselves relations to the environment. They are causally local to the individual. Many of them are internally accessible to the individual through reflection. Studying those states and their relations to one another can, for many of the purposes of systematic psychology, ignore the environmental relations. It might even turn out that systematic study of the relations between individuals and their environment is scientifically fruitless. I think that this is not the case; and I shall return to the point. But its being the case would not establish individualism (internalism), or undermine anti-individualism, as I understand those positions. It would not even show that the psychological states studied by scientific psychology were not anti-individualistically (“externally,” in *my* sense) individuated. Internalism, in *my* sense, concerns not the locus of the psychological states, or the best ways to study them, but whether being in them presupposes individual-environmental relations. It concerns whether the existence and nature of certain psychological kinds depends necessarily on the existence and nature of certain relations to specific kinds or situations in the environment.

To recur to a well-worn analogy: What it is to be a heart depends essentially on what the heart does in the context of the body. It pumps blood to other parts of the body. A chemically identical object that did not pump blood would not be a heart. Imagine that it had evolved to carry out an entirely different function in an organism entirely different from any animals with hearts. Similarly, parts' being certain valves and ventricles in the heart depend on their being parts of a heart, which in turn depends on the heart's functional and causal relations to the wider “environment” of the body. In this sense what it is to be a heart depends on relations between the heart and things outside the heart. But one could still study the “internal” physiology of the heart—states and structures that are purely “inside” the heart. Their being the structures they are in some cases depends on their functions and on the heart's relations to the rest of the body. But their mechanics and internal structural features need make no explicit reference to this wider “environment.”

So it seems that on one of Chomsky's understandings of ‘internalism’, internalism is the view that the study of belief and meaning concerns internal states of an individual. On my understanding of ‘internalism’, however, internalism is the view that all psychological states, including beliefs and understanding of meanings, are completely independent—in a metaphysical rather than causal sense—of any individual-environment relations for being what they are. This difference in construal of ‘internalism’ seems to me to explain

why Chomsky makes a number of points that are intended to support “internalism,” or to raise difficulties for “externalism,” but with which I agree. These points are completely compatible with anti-individualism (externalism) in my sense.

For example, Chomsky sometimes writes as if Putnam’s initial view that meaning is not “in the head” defines opposition to internalism (Chomsky 1995, p. 44). Meaning is abstract and hence not anywhere. But the psychological state of understanding a meaning is naturally seen as “in” the mind or brain. Nothing in anti-individualism requires rejecting this natural view. I accept it (Burge 1982).

Similarly, Putnam’s criticism of the approach to the traditional study of meaning and reference does not define anti-individualism or externalism, in my sense. With his original thought experiment, Putnam attempted to show that it was not both true (a) that understanding the meaning of a term involves being in a certain psychological state (or associating a concept with the term), and (b) that the meaning (or the concept) determines the reference of the word (Putnam 1975).⁴ Contrary to Putnam’s arguments, I believe that both of these principles are true for a wide range of terms and concepts. In particular, (b) is true for terms or concepts that are nonindexical—terms or concepts whose range of application is fixed (up to vagueness) in a way that is definite and stable across different occasions of use and does not become definite only on occasions of use. For example, I can say apriori that ‘aluminum’ applies, if to anything, to aluminum; and I can do so, understanding aluminum as a specific kind. My understanding is constant across different occasions of use and does not become definite only on occasions of use.⁵ Chomsky notes that Putnam’s critique of the conjunction of (a) and (b) does not bear on I-meaning. But the critique is not representative of anti-individualism. In fact, it is incompatible with, or at least uncongenial with, anti-individualism in my sense.⁶

Chomsky criticizes views that suppose that lexical items have no internal meaning, but do nothing other than “denote a semantic value external to the person” (p. 271). On such views, semantic value just is external denotation. I agree with Chomsky that such denotational views of meaning and belief, if seen as giving a complete account, are without psychological plausibility or interest. One can reasonably insist that there must be more to meaning than “external” denotation, and more to propositional attitudes than relations to directly referred-to denotations. This “something more” must be somehow represented—at least as a structured state of understanding it, or of having an attitude typed by it—in the mind or psychology of the individual. But whether the relevant “internal” psychological and semantic kinds are dependent for their meaning, content, or nature on relations to an environment beyond the individual is a question left completely open by such insistence.

A more subtle example is provided by Chomsky’s discussion of Albert and Bill’s looking at indistinguishable apples A (Albert’s apple) and B (Bill’s). He writes, “Shall we say that

the thoughts, visual images . . . are the same for Albert and Bill but ‘directed’ to different things? Or different for Albert and Bill, the external objects A, B being ‘part of’ the thoughts?” (p. 282). He goes on to identify the second view with “externalist” usage, and claims that for the science of human nature an “internalist” picture seems appropriate. This passage is understandable given Chomsky’s construal of internalism as a claim that relevant psychological states are internal to the mind/brain. But an anti-individualist (or externalist in my sense) could accept either view. In fact, I do accept the first view with minor qualifications, the one Chomsky favors.⁷ Differences between Albert and Bill’s minds go no further than different token applications of their common perceptual and conceptual content. But their having beliefs about apples at all—as well as about the particular apples that they are related to—is individually dependent on their bearing certain relations to an environment.

III

None of this is to say that there is no real disagreement. It may well be that Chomsky accepts internalism and rejects anti-individualism *in my senses*. It may well be that he believes that all mental/psychological states (as distinct from relations between psychological states and things people refer to) are what they are in complete independence of the nature of the environment, beyond the individual’s body. But I know of no forceful, specific arguments for this position, or against mine. I want to discuss some places where Chomsky’s “internalism” seems to be straining toward the position I reject.

One such place is Chomsky’s apparent scepticism about the Twin Earth thought experiments. As I have said, some of his scepticism seems to be directed against the idea that the thought experiments show that meaning or belief is not “in the head.” But sometimes he appears to oppose understanding natural kind words like ‘water’ in ways supportive of Putnam’s original cases. And he appears to want to utilize such opposition to oppose the results of the thought experiments. Sometimes it appears that he rejects the Twin Earth methodology altogether. He suggests that it depends on cases that are too strange, or too subject to variation with slight changes in context, to provide stable intuitive judgments. And he implies that the Twin Earth methodology provides no more than data about people’s commonsense beliefs about belief and meaning, with no significant implications for belief and meaning themselves.

It seems to me that Putnam discovered something important about how a significant range of “kind” words actually work, in many people’s idiolects. I think that it is true that water has turned out through chemical analysis to be constituted of H₂O molecules, and that anything that is not so composed is not water. Most of us who know the chemical facts and who use the word ‘water’, or counterparts in other languages or idiolects, would

speak as I have just spoken. And even those who do not know the chemical facts are open to finding that water has a “nature” that is not necessarily dictated by its superficial characteristics.⁸

This point can be hedged to allow for vagueness or borderline cases. But heavy water, often cited as a difficulty, is strictly a form of H₂O; and we seem to count it a special form of water as well. The point is not impugned by Chomsky’s tea examples, or indeed Putnam’s examples about the dirtiness of Lake Erie or about coffee (see Putnam 1990). We all recognize that tea, coffee, and Lake Erie are each mostly water. A body of liquid can be tea while also being water with small amounts of tea in it. Whether and in what contexts we are inclined to point to a cup and *call it* ‘tea’ as opposed to ‘water’ does have to do with contextual conversational factors, but has little to do with our beliefs about what is water and what is not water. Such cases are not decisive, or even clearly relevant, in my view, for determining the semantical application of the term ‘water’ in individual idiolects. Theories of the reference-fixing of the term or concept through examples must make allowances for impurities. The anti-individualist account of these terms and concepts can easily accommodate these matters.

What is more important is that the anti-individualistic force of the thought experiments does not depend on the assumption that water is H₂O. The water-type thought experiments can be understood this way: All of us recognize that for a large number of terms, including ‘water’, it is in principle possible for something to look or seem to fall under the term, but fail to do so. We can make mistakes, by our own lights. We recognize that any one of us could be fooled into thinking that something is water—even on moderate normal experience with the stuff—and later be convinced that it is not water. One might deny this with respect to some terms. But most people would agree that there are many terms in their idiolects for which this is true. Agreeing is fundamental to recognizing the independence of the world from our beliefs. The relevant terms include not just natural-kind, constitution-driven terms, like ‘water’ and ‘aluminum’, but most other terms for empirically experienced objects and properties. By hypothesis, we imagine something that does not fall under the term that would systematically fool one into thinking it does so. One would be fooled short of deeper investigation than one has actually undertaken. The possibility of such a thing has already been granted by most of us, with respect to a wide range of terms. This is the status of the hypothetical stuff XYZ.

The point that water is not XYZ does not depend on a prior assumption that water is H₂O. XYZ is introduced as something that is not water. As long as one takes it to be a possibility that something could fool one in some relatively systematic way—because what water is depends on more than its superficially experienced properties—one is committed to XYZ’s not being water. To argue that XYZ *is* water would be to misconstrue the thought experiment.

The rest of the thought experiment depends merely on the point that one could not think about, or refer to, water *as water* if one were in the same position with respect to XYZ as the protagonist is in with respect to water. That is, if one had no account of the nature of water that went beyond one's experiences, and if one had experiences only with XYZ and contact only with others who had experiences only with XYZ, then one could not specify or think of anything as water. I think that this point is apriori and, on sufficient reflection, self-evident.

This is the key reasoning underlying the natural-kind thought experiments. It does not depend on science fiction, Twin Earth, or anything particularly esoteric or hard to evaluate. The reasoning is, I think, decisive.

The key step in the water-type thought experiment depends only on one's acknowledging, with respect to a wide group of terms or concepts in one's own idiolect or system of thought, that the correct range of application is not completely fixed by the way the things look or the descriptive properties one knows to ascribe to the objects that one calls by those terms. One's knowledge and ordinary common experiences with the objects need not be sufficiently specific to fix the range of things that the terms or concepts in fact succeed in specifying. Or if one likes, one's present abilities to describe do not suffice to fix the range of application of one's reference with the term. Virtually anyone will recognize such a group of terms in his idiolect—a very large and varied group—given sufficient explanation of what is at issue. Actual usage will commonly be in accord with such acknowledgment. No individual's experience and knowledge, or behavioral or physical responses, can be expected to be sufficient to fix, by analytical or metaphysical necessity, the nature of many of the empirical objects and properties that he thinks and talks about. There is slack between what one knows and experiences, or more broadly how the world impinges on one and the nature of the things that one specifies. This fact, together with facts about particular conditions under which particular sorts of reference or intentionality are possible, drives the thought experiments. One need not rely on some general thesis about the constitution of water, or the way that a certain class of natural-kind terms works.⁹

Similar points apply to the social anti-individualist thought experiments. The arguments do not depend on general claims about a public language or about shared meaning. They depend on recognition of our ability to think and talk nonindexically about certain specific items, despite an understanding of the terms or concepts that is insufficient to specify the items in other terms, except insofar as the specification goes through a reliance on others.¹⁰

The use of thought experiments like the various Twin Earth thought experiments that I have proposed requires judgment as well as insight. Most such thought experiments are inevitably underdescribed. They do not function as proofs from self-evident principles. One must have exercised judgment about whether or not a detail, or an omission, or an

oddity, is crucial to the main point of the thought experiment. Some of the cases Chomsky discusses are simply hard cases. Hard cases do indeed sometimes make bad law. It is especially noteworthy that there is much that we do not know about language and concept development that makes judgment in developmental cases difficult in principle.

In some of the cases, however—Chomsky’s robin case, or his whale-fish case—the underdescription is crucial. Fuller description, properly targeted, will produce a much “easier” case. One needs to know more about the speaker’s usage and dispositions in order to settle the case. An individual’s simply belonging to a community in which many people know that whales are not fish does not itself ensure that the word-sound ‘fish’ in the individual’s idiolect does not apply to whales. But the case can be settled if more relevant features of the individual’s disposition and usage are included in the example. (See also my reply to Owens.)

The sofa case that I produced and that Chomsky discusses does not depend on an attempt to provide a general characterization of radical disagreement about the nature of empirically identifiable objects.¹¹ It depends on the possibility, in particular cases, of shared reference between the disputants; on a possible gap between what we know about the referent—or denotation, or nominatum—and what the nature of the referent in fact is; and on the impossibility of reference to certain kinds of objects (sofas) under certain conditions. I agree with Chomsky that we need deeper clarification of the fundamental notions. I think, however, that not all of the relevant issues are empirical. I believe that the claim that the cases I sketched are possible, as described, is very solid and stable.

IV

I would like now to discuss methodological issues. The thought experiments that I have proposed depend on “intuition”—reflective judgment—about cases. Chomsky expresses discomfort over the use of intuition. I do not think that his point that the cases use semantic terms, like ‘refers’ or ‘denotes’, that are partly technical carries any weight against the cases. The cases can be explained so as to connect with widely understood concepts about language–world relations. Insofar as Chomsky’s discomfort is grounded in his experience in theoretical linguistics, however, I find the discomfort understandable. Intuitive judgments about the grammar of a language or about what mentalistic structures underlie certain types of competence have often turned out to be badly mistaken. In many instances Chomsky is right in saying, “intuitive judgments are data, nothing more.”

But I think that if this slogan is generalized, it greatly overstates the case. Sometimes intuitive judgments about belief, meaning, or reference constitute knowledge—certainly warranted belief. The point is embedded even in Chomsky’s own methodology for

studying lexical items. He holds that in having a language, an individual can know that ‘chase’ in his idiolect entails ‘follow’, or that chasing entails following. The reflective judgment on which this knowledge rests need not be buttressed by some empirical explanatory theory. It is apriori for the individual who judges it rightly. The judgment constitutes knowledge, both about words and about a relation between chasing and following, that does not rely on empirical investigation. This is not to deny that it is desirable to fit the knowledge into a broader explanatory theory. It is not to say that such judgments are immune to correction. But it is to say that intuitive judgments can provide knowledge about the topic of the judgments, not just data for knowledge about the judge.

The most secure home for such judgments is mathematics and logic. It is clear that meaning and belief are normatively constrained by broadly logical and mathematical structures. I think that in complex ways they are also constitutively constrained by such structures. Some of these structures are accessible to reflective “intuitive” judgment. Often the reflection goes by way of complicated, dialectical, self-correcting processes. And it would be absurd to claim infallibility or even reliability for any and every off-the-cuff judgment. Finding the scope and limits of reflective judgments, as sources of objective knowledge, is itself a complex matter. In some cases, the dialectic leads to something like a systematic theory, and such theory may include empirical elements. But the structural aspects of belief and meaning that are accessible to reflection go well beyond what is codified in traditional logic and mathematics.

The thought experiments that I have developed are ways of exploring and clarifying, through intuitive judgments, structural aspects of reference, meaning, and belief that are accessible to intuitive reflection. I believe that intuitive judgments in this domain are more likely to be secure when they fix on particular cases. The cases instantiate more general, structural principles. But exactly what principles they instantiate requires reflection, comparison with other cases, invocation of background knowledge—sometimes including empirical knowledge—and theory development. Here too intuitive judgments are more than data for empirical theory about the individual doing the judging. They provide warranted belief, and sometimes knowledge about reference, meaning, and belief. Although such judgments often make use of broad, well-established empirical background knowledge, they characteristically have elements that are apriori warranted.

Some apriori warrant is *prima facie*, or *pro tanto*. Then the warranted judgments can be overturned. Some apriori warrant is not *prima facie*, but depends on having achieved a fully secure understanding of the principles underlying the cases. Such understanding is hard to come by. Reflection is a complex and delicate enterprise. In both *prima facie* and non-*prima-facie* cases, intuitive judgments are vulnerable to correction. The judgments that purport to be non-*prima-facie* apriori warranted can be corrected because they rest on misunderstanding and are not warranted at all, though they seemed to be. Sometimes the

correction can come only from further apriori reflection. Sometimes the correction is empirical. But vulnerability to correction, even empirical correction, does not itself prevent a *warrant* for an intuitive judgment from being apriori.

Kripke and Donnellan offered compelling cases for the possibility—in particular cases—of a name's referring to a specific object, or a person's referring with a name to a specific object, not perceptually present, even though the person could not fix the object with a description (Kripke 1972; Donnellan 1970, pp. 335–358). Their judgments provide knowledge. The judgments do not need support from further empirical investigation. Empirical investigation can help show where such possibilities are realized. But I believe that it cannot undermine the possibility that they illustrate, or even show that the possibility is never realized in our actual use of names. Intuitive judgments of the sort offered in these thought experiments are not infallible. But these have a certainty that is not seriously in danger of being overthrown. Empirical theory is not irrelevant to understanding them. But the examples capture how reference works in a way that is accessible to reflective judgment. Empirical theory, insofar as it does not change the subject, will in some measure have to conform to the thought experiments. The same can be said of some of the anti-individualist thought experiments.

Not all warranted reflective judgments in the thought experiments that I have discussed are apriori warranted. Many involve a mixture of empirical and apriori elements. Disentangling the two is an important philosophical project. But the empirical elements are broad, well-established points that are relatively uncontroversial.

It seems to me that we should be trying not only to find universal, innate components in our linguistic and cognitive structures and abilities. We should also investigate what components can be known apriori. And we should investigate what can be known or warranted about other matters through reliance on these cognitive structures.

The apriori component does not coincide with the innate component. There are innate components that can be known only empirically. There are innate structures whose use in acquiring knowledge about other matters provides only empirical warrant and knowledge. For example, many of the categories and transformational mechanisms in our perceptual systems are innate; the warrant for their application is always empirical. And some of the apriori knowable and apriori warranted elements in our cognitive make-up are not innate.

The apriori component need not coincide with what is necessarily present in any thinker or language speaker. Thus, it does not follow from the fact that a Martian speaker, or even some human speakers, need not have a certain linguistic or cognitive structure that we cannot know it through apriori reflection. It also does not follow that we cannot obtain apriori knowledge of other matters by relying on it. Even if certain mathematical concepts are not universal to all speakers or thinkers, those concepts can be used to provide apriori warranted knowledge of mathematical structures.

There have been many psychological studies in recent years to show that our intuitions about our grammar, motivation, forms of reasoning, and so on, are in certain cases very unreliable. This has led many to hold that reflection is not a source of knowledge except insofar as it is made redundant by empirical theory. I think that this is an overreaction, born of a simplistic conception of apriori intuitive reflection. In particular, one hears the recurrent claim—as if it were an exciting, revolutionary insight—that intuitions and intuitive reflection are fallible, and that it is easy to fall into groundless, entrenched prejudice while purporting to rely merely on intuitively obvious truths. These points are correct. But they are no surprise to traditional proponents of reflection as a source of apriori knowledge. No serious rationalist has ever regarded attempts at intuitive reflection as infallible. No serious rationalist has seen knowledge by reflection as easy to come by, except in a few cases. Such knowledge commonly requires skill, reflection on a variety of examples, dialectical or other inferential checking, and an openness to correction. Examples of misuse of reflection abound. Descartes himself emphasized this point repeatedly. For all that, I believe that more of our psychological and semantical structures are apriori accessible, through careful reflection, than is commonly thought. I believe that we are also sometimes apriori warranted in acquiring, sometimes unreflectively, knowledge of nonpsychological, nonsemantical matters, by making use of structures embedded in our minds and languages.

Not all of what we can know through reflective judgment need be a part of systematic science. For example, I think that some of what we know about reference may never fit into a systematic scientific study of (say) the sorts of causal relations that enter into reference. What we know about reference through reflection is either very particular and case-bound, or very general. The fine structure of the cases may not lead to a fruitful science. I think that, in itself, such a result would not make our knowledge even slightly less solid or less genuine. So I do not accept systematic empirical science as the only arbiter for all knowledge of language or of psychology. But we should, of course, develop systematic empirical science where we can.

V

I have been discussing ordinary reflective judgments that may produce knowledge without being elaborated in an empirical science. What of the status of cognitive psychology and empirical linguistics? Are these “internalist” in my sense? Elsewhere I have given two types of arguments that they are not. One is that they use psychological and semantical terms or concepts (*believes*, *perceives*, *means*, *refers to*) that are relevantly similar to those used in the thought experiments. So the terms and concepts of the relevant sciences and the psychological kinds that they pick out are subject to the same considerations. Since I

believe that these sciences must use relevantly similar concepts in order to solve the problems that they set themselves, I believe that the anti-individualism of these sciences is necessary and fundamental to their basic objectives.

The second type of argument consists in reflection on particular theories in these sciences. Numerous theories in psychology and semantics that use standard scientific methodology, and that produce results that are as solid and potentially as rich as those of any purely syntactic or phonological theory, are manifestly anti-individualist or noninternalist in my sense.¹²

Chomsky's discussion of individual theories is inevitably brief and general, and hence hard to evaluate. But at least some of the discussion seems to depend on the difference in his understanding of 'internalist' that I noted earlier. He points out that some studies of vision use, instead of real-world objects, tachistoscopic presentations that cause a subject to have a visual experience as of a rotating cube. He points out that the same investigations could have proceeded by stimulating the retina directly, or the optic nerve. Similarly, studies of hearing could go by way of direct stimulation of auditory receptors, rather than by way of loudspeakers.¹³

These points certainly suggest that these investigations of particular aspects of perception concern the processing of internal representations. But if one thought that they presented any difficulty at all to anti-individualism in my sense, one would be misinformed. Anti-individualism does not require that one study only *cases* of veridical perception. It does not stipulate that one study veridical perception only by fixing on cases of veridical perception: One can learn a lot about perception by simulating veridical cases. It does not even stipulate that one study veridical perception, although psychologists of perception normally do. One can study perceptual representation. But the representation of a rotating cube in cases where there is no actual rotating cube to be perceived, requires some ability to represent cubes and motion. Anti-individualism holds that this ability is not possible unless there are relations between the perceptual system (not necessarily the individual's system, but the system the individual shares with other members of the species, or perhaps other species) and objects in space. The meaning or nature of the representational states depends on these relations. Science need not study these relations, but they have to be in place for science to study representations of those kinds.

Similarly, in the discussion of Marr's theory of vision, Chomsky claims that only Marr's informal pattern fails to be 'internalist'. The serious theory is, according to Chomsky, 'internalist'. One of his arguments for this view consists in claiming that 'representation' in Marr's account is "not to be understood relationally, as 'representation of'" (Chomsky 1995, pp. 52–53). 'Representation' in Marr's theory is indeed not in general the relation of successful, extensional representation of. But this provides no support at all for internalism as I understand it. In the passages in Marr that Chomsky cites, 'representation'

means ‘representation as of’. On this usage, a representation would be, approximately, a percept. ‘Represent’ as a verb is followed by a grammatical object in an intensional context. Representations function to represent certain properties, kinds, or tropes *as* real. Representations are intentional items, items with aboutness properties. The anti-individualist argument applies to them, or at least to a subject’s having them.

So a lot of Chomsky’s discussion of Marr’s theory defends a view that I accept. It does not engage with my position. There are also, however, central aspects of Marr’s theory—and indeed virtually all serious theories of vision—that Chomsky fails to call attention to. The theory’s main objective is explicitly stated to be that of explaining how we visually determine the properties of actual objects in physical space that we in fact visually represent as being as they are. I shall quote Marr at some length:

The purpose of these representations [in early visual processing] 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, and I shall be quite careful to do this.

From an information-processing point of view, our primary purpose now is to define a representation of the image of reflectance changes on a surface that is suitable for detecting changes in the image’s geometrical organization that are due to changes in the reflectance of the surface itself or to changes in the surface’s orientation or distance from the viewer. . . .

Hence we can see in a general way what our representation should contain. It should include some type of “tokens” that can be derived reliably and repeatedly from images and to which can be assigned values of attributes like orientation, brightness, size . . . and position. . . . It is of critical importance that the tokens one obtains 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 structure of the surface will be meaningless. (Marr 1982, pp. 43–44)¹⁴

. . . the true heart of visual perception is the inference from the structure of an image about the structure of the real world outside. The theory of vision is exactly the theory of how to do this, and its central concern is with the physical constraints and assumptions that make this inference possible. (Ibid., p. 68)

These points cannot be passed off as informal patter, as Chomsky does. They are the most serious, explicit statement of the objective and method of the theory. The method is to individuate representations and constrain their processing in such a way that one explains how they are reliably caused by and how they reliably represent the physical properties in the environment beyond the perceptual system that they in fact reliably represent in normal circumstances.

Moreover, Marr’s theoretical constructions consistently *follow* his stated methodology. A central formal assumption in the construction of the account of the primal sketch—a

representation of simple edge patterns in the physical world—is the spatial coincidence assumption. This assumption holds that the coincidence of zero-crossings in the outputs of different-sized filters indicates the presence of a feature of the visual image that is due to a single physical phenomenon, such as a change in reflectance, illumination, depth, or surface orientation (Marr 1982, p. 70). The justification of this assumption makes further use of empirical assumptions about the nature of the physical world, which Marr sums up as “the constraint of spatial localization” (p. 68). Moreover, earlier, the choice of filters is motivated partly through considerations of what filter best simulates our detection of physically real edges (pp. 54–61). This is just one of many possible examples. Detailed acquaintance with the theory will not allow one to dismiss Marr’s assumptions about the physical world, and about our visual detection of features of it, as informal patter. They are explicitly constitutive of the heart of the theory.

Chomsky has criticized my talk of a perceptual system’s solving a problem or having a purpose (Chomsky 1995, p. 55). He is surely right in warning against anthropomorphizing the perceptual system, or interpreting this talk in a way that would ascribe intent or deliberation to it. But some reference to a system’s provision of physical information about the environment through states with intentional (aboutness) properties is inevitable in any scientific account of perception. In this respect, psychology differs from chemistry or physics. Explanation of a system’s specific abilities in representing or registering¹⁵ or perceptually presenting features of the world is the primary traditional task of the psychological theory of perception. One can imagine psychological theories that totally prescind from this task. But I cannot see that such theories would be of any interest, at least as theories of perception. They would have changed the subject.¹⁶

Chomsky’s view that misperception by an animal or person is merely a matter of what people, observing the perceptual system, “decide to” call a “misperception” (Chomsky 1995, p. 53) does not accord with the actual explanatory practice of the psychology of perception. That practice takes the explanation of illusions and misperceptions as one of its tasks (Wandell 1995, chapter 11). Illusions are facts that constitute part of the subject matter of vision. They indicate real tendencies within the visual system, or in its relation to the world, that make it deviate from the norm of presenting veridical visual information about the physical world.

Another area where explicitly relational elements enter into empirical theorizing is the formal semantics of natural language. In my view, this area is less advanced than the psychology of perception. But I see no reason why it should not be an area of fruitful systematic scientific investigation. I see no reason why reference, or a technical analogue, as a relation between linguistic representations and real aspects of the world, should not be partly systematizable in a formal account of the truth-conditions of sentences and utterances.

I agree with Chomsky's criticism of referential semantical theories that purport to do away with, or ignore, I-meaning—or with the way individuals understand their terms—in favor of a purely denotational or referential semantics. But I cannot sympathize with his approving quotation of Hobbes's claim that it is “manifest” that names are signs not of things but of our cogitations. I think that the negative part of Hobbes's view is manifestly mistaken. It is as mistaken as applied to our ordinary commonsense use of language as it is to scientific language. The idea that we name and talk only about our cogitations seems to me beyond serious discussion.

As I noted earlier, I think that we can know things about referential or denotational relations even if we cannot systematize them scientifically. Scientific systematization is dizzyingly difficult and complex. I do not, however, think that the sorts of considerations that Chomsky raises provide ground for believing that the enterprise is fruitless. The points about the differences in ways nominal expressions relate to the world ('Joe Sixpack' as opposed to 'Julius Caesar') are part of our intuitive self-understanding and have to be accounted for in a formal semantics that purports to systematize language–world relations. The Joe Sixpack example is not hard to provide at least an approximate gloss on.

It may well be that for some purposes, semantic values can be taken to be shorthand constructs with no simple real-world correlates. But insofar as it does this, semantic theory should distinguish between the semantics of this sort of name and the semantics of ordinary names of people. For such a distinction is part of our ordinary linguistic understanding and use. It should explore the relations between such constructs and those things in the world that make talk about the constructs capable, often as a kind of shorthand, of contributing to true statements.

I see no reason to believe that semantic theory should collapse into the idealist and unsupported philosophical ideology that naming is in general nothing more than a kind of “worldmaking.” What is true is that a good linguistic theory must say something about the understanding or cognitive value of names as well as something about their reference.

A successful theory of reference cannot reasonably confine itself to the austere ontology of physics. It will allow objects like colors, tables, and symphonies. These are objects that we believe to be real even if no science studies them *per se*. Insofar as Joe Sixpack is real, “he” is not a person, or a single thing. Language–world relations are varied and complex, but variety and complexity are obstacles faced by all sciences.

Taking account of language–world relations is part of the way semantics is actually practiced. I see no reason to think that there is anything scientifically wrong or fruitless in studying language–world relations, or with taking them to be part of the formal structures elaborated in semantical theory. Like aspects of the theory of perception, this aspect of semantics is not internalist, even in Chomsky's broad sense of internalist. No serious

scientific objection has been entered against this practice. It may well be, of course, that systematic science will be more successful in studying formal structures into which referential items enter than in studying the ways such items connect to the world (i.e., in studying the reference relation itself).

Quite apart from the fate of systematic formal theories of reference or denotation, theories that make use of semantic features of lexical items are still subject to anti-individualist strictures. Even if these semantic properties of lexical items are “determined by the ways they are constituted, with a rich innate contribution,” the constituent elements depend for their meaning on relations between the individual and an environment.¹⁷ The marking of the gender of a noun, the distinction between persons or animals and nonliving things, the indication that a term ‘chase’ entails ‘follow’ (where ‘follow’ is understood in the normal way in terms of moving in the same direction as, but behind)—all presuppose mind–world relations between the cognitive/linguistic system and the world beyond the system. This aspect of semantics is anti-individualist in my sense even if it is ‘internalist’ in Chomsky’s. I believe that the study of lexical meaning is no more advanced and no more clearly scientific than the study of truth-conditions and reference. But anything that we recognize as a study of the meanings of a wide range of ordinary words or lexical items will have substantial anti-individualist presuppositions.

There is no need to defend generalized internalism in either sense in order to maintain Chomsky’s brilliant insights and methods in the study of the language faculty. That faculty involves a rich set of specific unconscious structures and rules, many of which are innate. The psychological states typed by phonological and syntactic structures are clearly internalist in Chomsky’s sense.¹⁸ But scientific enterprises that study language and mind and that are not internalist in Chomsky’s sense are already well launched. Many of these enterprises feed on Chomsky’s insights in syntactic theory. Reinterpreting these enterprises to fit an internalist paradigm, or counting them, on general or anecdotal grounds, unscientific or bound to fail, will not carry conviction—any more than philosophically motivated criticisms of the manifestly fruitful methods and results of Chomskian syntax have carried conviction. Formal semantic theory, which includes formal theories of truth-conditions, seems to have some promise as a scientific enterprise. I have no doubt that some aspects of what we know about reference will not be systematized into a science. Some knowledge is not systematic or scientific. How much of a theory of reference stabilizes as a science remains to be seen. The upshot will not be determined by philosophical views, but by the rough and tumble of linguistic scientific investigation.

My primary point, however, has not been to defend scientific studies of the relation between individuals and their environment. It is to claim that even actual scientific enterprises within cognitive psychology or semantics that do not study individual–world relations, and are internalist in Chomsky’s sense, are not internalist in mine. Many of their

basic psychological kinds are dependent for their natures on relations to an environment. The relations are presupposed, rather than made a topic for scientific study.

Some knowledge derives not from empirical science, but from reflection. Perhaps some of this knowledge will never be systematized in a science. But some of it is presupposed by and even frames scientific reasoning. Where scientific reasoning comes to grips with human nature, it will have to deal with systems that have intentional or aboutness properties. For human nature is partly constituted by our rational, cognitive capacities; and these capacities are fundamentally intentional or representational. Our rational and cognitive capacities represent a world that we can make mistakes about. These facts, together with apriori-knowable conditions on the possibility of such intentional representation, make it inevitable that many of our intentional representations are constitutively dependent on certain nonintentional relations between our cognitive systems and a world beyond them.

Notes

1. In "Language and Nature" (1995), p. 31, Chomsky quotes my description of the eliminativist view in a way that might suggest to someone who does not know my work that I take the view seriously. In fact, I think that the existence of mental states and events is as epistemically solid as the existence of rocks and trees.
2. See my (1989).
3. This understanding of 'internalism' seems to motivate much of Chomsky's (1995). For example, he writes in introducing 'internalism': "Internalist naturalistic inquiry seeks to understand the internal states of an organism" (27; see also p. 46). He gives as examples of "non-internalist studies of humans" the study of phases in the oxygen-to-carbon-dioxide cycle or the study of gene transmission, or of individuals as participants in associations and communities (28).
4. See my (1982) for my criticism of Putnam on these matters. Putnam has accepted my central criticism of the claims of his original, brilliant paper. See the introduction to Putnam (1996). My criticism about indexicality undermines his rejection of the conjunction of the two central principles that he says in "The Meaning of 'Meaning'" cannot both be held. These points are laid out in "Other Bodies." For later anti-individualist work by Putnam that is more congenial to my view, see "Computational Psychology and Interpretation Theory" in Putnam (1983).
5. Chomsky may doubt that there are any such terms when he holds that "reference" or language-world relations is confined to something a person establishes only on particular occasions of use. I see no ground for this view. It seems even less plausible as applied to thought than as applied to terms.
6. See note 4 and Putnam (1975). Putnam now rejects his own earlier account of the thought experiments.
7. I discuss this very case, down to the example of apples, at the beginning of my (1982).
8. A systematic empirical study that supports this view about natural kinds has been done by Susan Gelman. See Gelman, Coley, and Gottfried (1994).
9. Some have suggested that the mental content might leave an open parameter for the referent of the mental representation. This method seems to me to have little relevance to the way we actually think. Our thoughts commonly have specific commitments to specific referents (kinds, properties, relations, especially). See further discussion of this point in my replies to Owens, Peacocke, Loar, and Donnellan.
10. I have discussed these points in greater detail in my (1989), as well as in my (1979).
11. See my (1986b), pp. 697–720.

12. For an example of such a theory, beyond the visual perceptual cases I have already discussed in “Individualism and Psychology” (1986), pp. 3–45, see Hutchins (1995). For discussion of a range of other such theories, primarily in psychology, see McClamrock (1995); Silverberg (1998); and Wilson (1995).

13. See Chomsky (1995), p. 52. Chomsky claims that these theories apply exactly as well to a brain in a vat as to an individual perceiving actual objects. I think that this is true, but only because such a brain has a perceptual system that has representational abilities typed by relations that other brains have borne to their environments. See my reply to Loar.

14. I indicated the anti-individualist caste of Marr’s theory at some length in my “Individualism and Psychology” (1986a). My “official” argument for the anti-individualist caste of the theory was unnecessarily complicated and blurred the distinction between failure of local supervenience and anti-individualism. But the discussion of the way Marr’s theory works made the basic point clear. The science associates the content of perceptual representations with their responding successfully to features in the environment according to certain formation principles, which also have their content by virtue of normally corresponding to facts in the environment. (See “Descartes and Anti-Individualism: Reply to Normore,” this vol., note 13.) A number of discussions have specifically doubted the basics of my account of the theory. I believe that all of these, including some that Chomsky cites approvingly, either manifestly fail to understand the theory or fail to understand the philosophical issues that my account turns on. There is also a large literature which accords with my account of the theory and its successors. For a recent example, see Georges Rey, “Chomsky, Intentionality, and a CRTT” (forthcoming), section V(ii). The response to Marr’s work in the psychological literature has not missed the fundamental role of features of the actual physical world in providing constraints on the theory’s attribution of the nature of perceptual content and processing procedures. See Grimson (1981), e.g., chapters 1, 2, and 9.4.2; Bruce and Green (1985, 2000), preface and passim; and Ullman (1996), passim. Moreover, the anti-individualist character of the empirical psychology of vision is quite general and in no way is confined to Marr’s theory. For other work that shows, very explicitly, anti-individualism in the explanatory methods and actual theory of vision, see Shepard (1984), Ullman (1979), Spelke (1990), and Palmer (1999).

15. This is Chomsky’s word. See Chomsky (1995), p. 52.

16. These points are illustrated in Wandell (1995), and in Ullman (1996). For example, in chapter 2 Wandell discusses in detail measurable responses of the retinal image to measurable properties of light. In chapter 3, he tests forms of representation against the viewing angle of the actual light and proposes a theory of how the cone mosaic encodes the high spatial frequency patterns created by visual interferometers, and how the spatial arrangement of the cones provides information about the spatial distribution of light. In chapter 4, he discusses what information about spectral power distribution in actual light is encoded when rods initiate vision. When the discussion turns to visual representation (as opposed to neural registration) in chapter 5, Wandell uses the discussion of the relation between actual light and neural registration in preceding chapters as a constraint on the nature and mechanism of visual representation. Chapter 10 includes a detailed account of the construction of representations of depth from representations of motion that—like the earlier chapters that I have glossed—is not purely internalist even in Chomsky’s sense.

17. It is understood throughout that such relations can be indirect. The relations relevant to the individual’s innate structures might go through the individual’s progenitors and their relations to the environment—relations that played a role in giving the innate structures their intentional content and function.

18. I believe that some elements of phonology and syntax are not internalist in my sense, but I shall not discuss these points here.

References

- Bruce, Vicki and Patrick Green. 1985, 2000. *Visual Perception: Physiology, Psychology, and Ecology*. Hillsdale, NJ: Lawrence Erlbaum.
- Burge, Tyler. 1979. Individualism and the Mental. *Midwest Studies in Philosophy* 4: 73–121.
- . 1982. Other Bodies. In *Thought and Object*, A. Woodfield (ed.). Oxford: Oxford University Press.
- . 1986a. Individualism and Psychology. *Philosophical Review* 95: 3–45.

- . 1986b. Intellectual Norms and Foundations of Mind. *Journal of Philosophy* 83(12): 697–720.
- . 1989. Wherein Is Language Social? In *Reflections on Chomsky*, A. George (ed.). Oxford: Blackwell.
- Chomsky, Noam. 1995. Language and Nature. *Mind* 104(413): 1–61.
- Donnellan, Keith. 1970. Proper Names and Identifying Descriptions. *Synthese* 21: 335–358.
- Gelman, S. A., J. D. Coley, and G. M. Gottfried. 1994. Essentialist Beliefs in Children: The Acquisition of Concepts and Theories. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, Hirschfeld and Gelman (eds.). Cambridge: Cambridge University Press.
- Grimson, W. E. L. 1981. *From Image to Surfaces*. Cambridge, MA: MIT Press.
- Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kripke, Saul. 1972. Naming and Necessity. In *Semantics of Natural Language*, Davidson and Harman (eds.), pp. 253–355. Dordrecht: Reidel.
- Marr, David. 1982. *Vision*. San Francisco: W. H. Freeman.
- McClamrock, Ron. 1995. *Existential Cognition: Computational Minds in the World*. Chicago: University of Chicago Press.
- Palmer, Stephen. 1999. *Vision Science*. Cambridge, MA: MIT Press.
- Putnam, Hilary. 1975. The Meaning of “Meaning”. In *Mind, Language, and Reality: Philosophical Papers*, volume II. Cambridge: Cambridge University Press.
- . 1983. Computational Psychology and Interpretation Theory. In *Collected Papers*, volume III. Cambridge: Cambridge University Press.
- . 1990. Is Water Necessarily H₂O? In *Realism with a Human Face*, James Conant (ed.). Cambridge, MA: Harvard University Press.
- . 1996. Introduction. In *The Twin Earth Chronicles: Twenty Years of Reflection on Hilary Putnam’s “The Meaning of ‘Meaning’”*, Andrew Pessin and Sanford Goldberg (eds.). New York: M. E. Sharpe.
- Rey, Georges. Forthcoming. Chomsky, Intentionality and a CRTT. In *Chomsky and His Critics*, Antony and Hornstein (eds.). Oxford: Blackwell.
- Shepard, Roger N. 1984. Ecological Constraints on Internal Representation: Resonant Kinematics of Perceiving, Imagining, Thinking, and Dreaming. *Psychological Review* 91: 417–447.
- Silverberg, Arnold. 1998. Semantic Externalism: A Response to Chomsky. *Protosociology* 11: 216–244.
- Spelke, Elizabeth. 1990. Principles of Object Perception. *Cognitive Science* 14: 29–56.
- Ullman, Shimon. 1979. *The Interpretation of Visual Motion*. Cambridge, MA: MIT Press.
- . 1996. *High-Level Vision*. Cambridge, MA: MIT Press.
- Wandell, Brian A. 1995. *Foundations of Vision*. Sunderland: Sinauer.
- Wilson, Robert. 1995. *Cartesian Psychology and Physical Minds: Individualism and the Sciences of the Mind*. Cambridge: Cambridge University Press.