ROBERT STALNAKER

MENTAL CONTENT AND LINGUISTIC FORM

Many philosophers have argued or assumed that language is, in some way or other, essential to thought. Philosophical explanations of the nature of intentional mental states, psychological explanations of the mechanisms of mental representation and semantic explanations of the attribution of content frequently rely in various ways on an analogy between linguistic and mental representation. Some of these philosophers explain thought as an internalization of speech, or as a disposition to speak, with the content of the thought being derived from the content of the speech act that is internalized, or in which the disposition is displayed. Some describe the internal states that constitute thoughts in terms of the compositional structure with which semantic theory interprets complex expressions. In previous work I have resisted projects of this kind, arguing for a conception of content that has no semantic structure, and for an explanation of the intentionality of thought that makes no reference to language or speech.[1] But it must be granted that lingusitic structure and linguistic practice are intimately involved in our mental life. Even if there can be thinkers who are not speakers or interpreters of speech, thinkers who do not in any sense talk to themselves or store their thoughts in words, it is clear that such thinkers would be very different, both in what they think about and how they do it, from us. In this paper I want to try to come to terms with what I have called the linguistic picture — the cluster of metaphors, analogies and strategies that use language in some way to help account for thought. I will try to separate out some of the different ways that mental content may be related to linguistic form, and some of the different ways that linguistic practices and objects may be involved in thought. Although I will make some concessions to the linguistic picture, I will continue to defend the course-grained conception of informational content, arguing that it is adequate to explain the various intimate relations between language and thought, and that it in fact helps to clarify those relations.

Philosophical Studies **58**: 129—146, 1990.
© 1990 *Kluwer Academic Publishers. Printed in the Netherlands.*

I am going to begin with an abstract semantical problem about belief attribution, although my main concern will be to use this problem to raise and sharpen some more substantive questions about belief: both philosophical questions about the nature of representational states and psychological questions about how such states are in fact realized. The semantical problem is to explain how the meaning or content of a statement attributing a belief is determined as a function of the meanings or other semantical values of its parts. The basic semantic structure of a statement of the form *x believes that P* is straightforward: it says that a relation (expressed by "believes") holds between the individual denoted by the subject term and whatever it is that is denoted by the sentential clause. The problem is to say, first what sort of thing it is that is denoted by the clause, and second, how a thing of this sort is determined as a function of the constituents of the clause. The semantical problem is concerned solely with the compositional structure of such sentences. All that needs to be said about the meaning of the simple expression "believe" is that it expresses a relation between a believer and an object of belief. But while a solution to the semantical problem can be stated without making any claim about the nature of this relation, the evaluation of competing hypotheses about what the object of belief is will require us to consider more substantive questions about belief.

It might seem that if our general semantic theory — our theory of constrcutions other than those involving discourse — is sound, then belief constructions — at least simple sentences of the form *x believes that P* — ought to raise no new problems. For if our semantic theory is compositional, it must be that the semantic value of a sentential clause is a function of the semantic value of the sentence from which it is derived. It seems natural to assume that the clause "that alligators are irritable" *denotes* what the sentence "alligators are irritable" *expresses*. So if our semantic theory tells us what "alligators are irritable" says, (or at least how what it says is a function of the meanings of its parts) then without any further work it will tell us what it is that someone believes when she believes that alligators are irritable. I think this is roughly right, but what it implies is not that belief constructions are unproblematic, but rather that any problems with belief constructions are symptoms of a much more general problem in the foundation of

semantics. And it is clear that there are problems with belief constructions — prima facie conflicts between what the phenomena suggest about how objects of belief must be individuated and what a truth-conditional semantic theory standardly says about the semantic values of sentences. Truth-conditional semantics (whether developed in the possible worlds framework or in some other way) takes the primary task of semantics to be to explain how the truth-conditions of a statement are determined. According to such a theory, statements with the same truth conditions — sentences that are true and false together under all the same conditions — say the same thing, although they may do it in different ways. But sentential clauses made from sentences that the semantics says are necessarily equivalent (and thus have the same truth conditions) often seem clearly to denote different objects of belief. It seems clear, for example, that one may believe that seventeen is a prime number without believing that the sum of the angles of a Euclidean triangles is equal to a straight angle. But these two propositions are both necessarity true, and so true under the same conditions.

This, then, is the semantic problem: to find an object of belief finely enough individuated to account for the phenomena of belief attribution, and to explain how such objects are determined as a function of the constituents of the sentential clauses that refer to them.

One might respond to the problem by holding to the thesis that objects of belief are individuated by truth-conditions and changing one's account of the truth-conditions of the problematic examples. Or one might deny that the referent of a *that* clause is a function of the meaning of the constituent parts.[2] But the strategy I want to explore is the most straightforward one: it takes at face value what the examples of belief attribution seem to show — that the object of belief is an entity individuated more finely than by its truth conditions, and it accepts the apparent consequence of this — that our semantics must attribute more fine-grained semantic values to sentences generally. The task then is to say just what sort of fine-grained objects are the right ones, and to see that one's semantic theory explains how such values are determined. There are many ways to carry out this strategy,[3] but the most promising way is to build some kind of semantic structure into the object of speech and thought. According to the standard truth-conditional semantic theory, the semantic structure of a sentence or sentential clause

is a part of the means by which the semantic value of the sentence or clause is determined, but not part of the value itself. The structure provides a recipe for determining content as a function of the values of the parts of the sentences, but the values of the parts are not constituents of content. The same content might be determined in quite different ways, as the value of different functions applied to different arguments. But while this is the way semantic theories standardly treat content, such theories do have the resources to define abstract objects that reflect the semantic structure of sentences — objects that represent the recipes themselves rather than the end results of following them. One might take the values of sentences and sentential clauses to be the structured meanings by which the truth-conditions, or informational content, of those sentences or clauses are determined. According to this approach, these structures will be the senses of sentences, and the referents of sentential clauses. They will thus be be the objects of both speech and thought.[4]

There are questions about the details of the development of this kind of theory. Should the structured meanings include Russellian propositions with individuals as constituents, or should they be structures containing only intensional objects such as senses or modes of presentation? At what level of abstraction should structure be analyzed? Might sentences with different superficial structures have the same structured meaning? There are some technical problems to be overcome — for example, avoiding the threat of semantic paradox — and there will remain some prima facie conflicts between a structured meaning account of the objects of belief and the phenomena of belief attribution. Sentences with the same structures and synonymous constituents may in some contexts seem to express different objects of belief. ("All woodchucks are groundhogs" and "All groundhogs are groundhogs," to take a familiar example) But I will assume that a theory of this kind that is formally adequate and more or less true to the phenomena can be developed. The question I want to consider takes us beyond the semantic problem — the problem of the structure of sentences attributing belief — to the substantive problem about the nature of the belief relation: what are the facts about the believer and his relation to his enviornment that make a belief attribution correct? If semantic structure is essential to the abstract object the believer is said to be related

to, how must that structure be involved in the state of the believer that constitutes his having the belief in order for the attribution to be correct?

However a fine-grained content is defined, it will be something that determines truth-conditional content, or what I will call *informational content.* Belief states, whatever they are, *could* be described simply in terms of their informational content. If x believes that P, then x is in a state that has the informational content determined by the referent of "that P." But according to the fine-grained strategy, a belief attribution says more than this, since even when P and Q determine the same informational content, it may be true that x believes that P, and false that she believes that Q. The question is, what is is about the believer's state — a state that has the informational content determined by Q (as well as by the equivalent P) — that makes it a belief that P, rather than a belief that Q.

To help sharpen and clarify this question, I will review a familiar, and I think plausible, answer to the prior question, what is it about a belief state that accounts for its informational content? States of organisms, according to this answer, carry information when there exists a pattern of counterfactual dependencies between those states and corresponding states of the envirnoment. If x is in a state caused by the fact that P, and would not have been in that state if it had not been that P, then that state of x carries the information the P. If x is capable of being in a range of alternative states that tend to vary systematically with variations in the environment, then those states will be information carrying states. While it may be controversial whether one can explain the content of intentional states entirely in such terms, it seems reasonable to assume that it is at least a necessary condition for a state's being a belief state that it tend to carry information in this sense. Further, it seems reasonable to assume that if a person believes that P then the relevant belief state tends to carry the information that P (even if carrying the information that P is not sufficient to make it a belief that P).

Representational states and systems can carry misinformation as well as information in the strict sense. We need an account of informational content that is neutral as to whether the world is as it is represented to be. But misrepresentation must be understood as a deviation from

a norm. It is reasonable to assume that representational states are *normally* correct — that they are states that *tend* to represent things as they are. Given an appropriate conception of normal conditions, we can explain representation generally in terms of information: a state represents the world as being such that *P*, and so is a state with informational content *P*, if and only if under normal conditions it would carry the information that *P*. Normal condition will include both conditions on the enviornment and conditions on the internal functioning of the representational mechanisms. The question of what the specific relevant normal conditions are is an empirical question, and answering it is part of the task of explaining the capacities and limitations of believers and other organisms that represent.

Information and representation in this sense are necessary for belief, but obviously not sufficient. For an information bearing state to be a belief state, the information must be available to play a role in determining how the believer acts to satisfy its wants and needs. If there are systematic dependencies between the state of an organism's environment and the way it is disposed to act, then the behavioral dispositions themselves will be representational states — states that tend to carry information about the enviornment. Beliefs, whatever else they must be, are presumably states of this kind.

Now let me return to the question about what kind of claim a belief attribution is making. The semantical hypothesis we are considering is that a belief attribution of the from *x believes that P* relates a believer to an object that has semantic form as well as informational content. The question is, what does such a belief attribution say about the believer's state beyond the claim that it has the informational content determined by the object of belief. Here is one kind of answer: just as a sentential clause has both informational content and semantic form, so an internal representational state of a believer will have both content and form. The information carried by such a state will be carried in some particular form, and, according to this answer, it is plausible to suppose that the form in which the information is stored is a linguistic form. Suppose that the way we represent the world is by storing sentences or sentence-like structures of a mental language — mentalese, or the Language of Thought. Then these representations will have the kind of structure that is repesented by a structured meaning or a

Russellian proposition. The claim that a belief attribution *x believes that P* makes, according to this suggestion, is that the way the believer internally represents the informational content expressed by *that P* is the same (at the appropriate level of abstraction) as the way it is represented in the clause *that P*.

It is important to recognize that the suggestion being made is not just a claim about what is going on in the believer; it is a claim about what a belief attribution says about what is going on in the believer. That is, the claim is not just that belief states take the form of internal linguistic representations; it is the claim that when we attribute specific beliefs, we say something about the specific semantic structure of the sentences of the internal language that encode those beliefs. According to this suggestion, if I say that *x* believes that *P*, my claim will be false if the form in which the informational content of *that P* is stored is relevantly different from the form of the clause "that *P*." I think this suggestion makes a belief attribution carry more theoretical weight than it is plausible to assume that it carries. If it were correct, belief attributions would be far more speculative, and believers would be far less authoritative about their beliefs, than they seem to be. While theoretical and experimental developments in cognitive psychology may someday convince me that I store my beliefs in a form that is structurally similar to the form in which they are expressed and described in English, I don't think that my ordinary belief attributions commit me to thinking that they will. Consider an example: Angus believes that Edinburgh is closer to Liverpool than it is to London, or so it seems. He acts as if he believes this; for example, when he wanted to go from Edinburgh either to London or to Liverpool, he didn't care which, and wanted to go as few miles as possible, he went to Liverpool. Furthermore, his linguistic behavior seems to support the hypothesis that he has this belief: he says, sincerely, and without hesitation, "yes" when asked, "do you believe that Edinburgh is closer to Liverpool than it is to London?" Angus thinks he has this belief. But does he, or do we, have reason to believe that the form in which this geographical information is stored in Angus is appropriately similar to the semantic structure of the sentence "Edinburgh is closer to Liverpool than it is to London"? Suppose the information is represented in Angus in a mental map. When planning his trips, or anwering questions, he consults his map, visualizes it

perhaps, and sees that London is farther from Edinburgh than Edinburgh is from Liverpool. Or suppose the information is stored in a quite different linguistic form — in a mental machine language whose semantic structure is quite different from English, although capable of representing the same information. Would this mean that Angus, despite what he thinks, does not really believe that Edinburgh is closer to Liverpool than it is to London? Surely not.

Those who want to explain beliefs as stored sentences of a mental language usually distinguish implicit beliefs from explicit or core beliefs. It is acknowledged that there is lots of information implicit in the believer and accessible if needed, that it is not plausible to assume is written down: unconsidered instantiations of universal generalizations believed, or propositions too trivial to notice such as, for example, that 4652 is an even number, or that no aardvark weighs more than an aircraft carrier. Such implicit beliefs are explained as dispositions to form explicit representations upon considering them.[5] Exploiting this distinction, we might say that even if Angus's representation of the informational content that Edinburgh is closer to Liverpool than to London does not *initially* have the right form, perhaps he would come to represent it in this form if it were presented it to him in this way, and so perhaps he at least implicitly believes this all along. But this too is speculation. Even when Angus is asked the question, in English, and gives his answer, there is no reason to think a sentence of the language of thought with just this structure needs to be stored. Angus does need to process the English sentence, to represent the fact that it has the semantic structure that it has. But this metalinguistic information is different from the geographical information that Angus may continue to store only in some quite different way.

There may be good empirical reason to believe that we store at least some information as sentences of an internal language. If one assumes a conception of language that is sufficiently broad and flexible, one might argue that anyone with the kind of sophisticated representational capacities that human beings have *must* store information in linguistic form. But I don't think it is plausible to believe that our ordinary belief attributions make claims about the specific structure of these representations. If belief is a relation to a fine-grained structured meaning, it is still not the structure of internal representations that makes them true.

I want now to consider a contrasting alternative account of what a belief attribution is saying about the semantic structure of the informational content of a belief. Perhaps the claim is not about how the information is stored or represented internally, but about how the believer is disposed to express the belief. Whatever is going on inside the head of the believer, if he is disposed to assert or assent to a sentence with a certain semantic structure, then it will be correct, according to this suggestion, to attribute to him a belief with this structure. If a person has a belief with the informational content *that P*, but sincerely dissents from sentences with the semantic structure of the sentence *P*, then, according to this suggestion, it will be false that she believes that *P*.

I think this suggestion about the role of semantic structure in the attribution of belief is more promising but it does require a closer conceptual connection between speech and thought than some will find plausible. The suggestion implies that there cannot be thinkers who are not speakers, at least not believers who are not speakers. For if someone or something is not disposed to express or communicate the information it represents in any form at all, then this proposal will imply that all attributions of belief to it will be false.

Not everyone will find this consequence unpalatable. Donald Davidson, for example, has defended a conception of belief that connects it essentially to speech. He makes such claims as that "making detailed sense of a person's intentions and beliefs cannot be independent of making sense of his utterances."[6] and that "we have the concept of belief only from the role of belief in the interpretation of language, for as a private attitude it is not intelligible except as an adjustment to the public norm provided by language."[7] Part of Davidson's motivation seems to be verificationist: he says, "we have no good idea how to set about authenticating the existence of such attitudes when communication is not possible."[8] But he also seems to be concerned with our problem: he suggests that we need speech and the interpretation of speech to make the kind of fine discriminations between intentions and beliefs that we seem to make.

While the conception of belief required by this response to the problem does have some independent support, it also has some facts to explain away. We do attribute beliefs to nonhuman animals and pre-

linguistic children, and it does not seem difficult to imagine cases of highly sophisticated but totally uncommunicative agents and inquirers. Suppose there were a species of solitary martians who go through life never seeing another creature of their kind. (They hatch from eggs long after their single parents have died.) Suppose that although these creature have no need or opportunity to communicate, they are highly intelligent creatures who observe, calculate, experiment, and formulate hypotheses about their environment, history and nature, and about their own thought. Perhaps they think in an internal mental language, but because of their prodigious memories and great powers of concentration, they have no need to externalize their thought — they make no pencil and paper calculations and keep no written records. If all of the alleged thoughts of such creatures remain internal, must we say that they have no beliefs? If we did attribute beliefs to them, would we have to understand our attributions in a highly counterfactual way — as claims about what they would be disposed to say if they had the resources and motivation to communicate? But of course how they expressed their beliefs if they did would depend on the particular means of expression that they are counterfactually assumed to acquire.

Even if we forget such counterfactual martians and restrict ourselves to the actual paradigms of thinkers — thinkers who are also speakers — we have to recognize some mental states that thinkers are not in a position to express. We make tacit presuppositions, have unconscious beliefs, and take things for granted without noticing that we are doing so. The correctness of the attribution of attitudes of this kind cannot be dependent on the means that we would use to express the information that those attitudes are carrying. Of course it might be that only certain attitude attributions make a claim about form as well as informational content. One might hold that attributions of conscious belief are correct only if the believer is disposed to assert or assent to a statement of the same form as the attribution clause, while also holding that the correctness of attributions of tacit presupposition depend only on the informational content of the attribution clause. But this would be plausible only if the original problem — the problem that motivated us to look for a more fine-grained object of belief — was a problem only for conscious belief. Our original semantic problem was that there seemed to be clear of belief that P without belief that Q, where P and

Q had the same informational content. But just as one can *believe* that seventeen is prime without *believing* that the sum of the angles of a Euclidean triangle is equal to a straight angle, one might also take the first for granted, or presuppose it in one's mathematical practice, without presupposing or taking for granted the second. This seems to imply that the root of the problem — the apparent fact that it is possible to believe something without believing everything equivalent to it — is not to be found in the fact that we are disposed to express our beliefs. This should be clear from the case of the solitary martian. Whether we say that it has beliefs or not, it is difficult to deny that it has some kind of intentional, information-bearing states. Call them schmeliefs instead of beliefs. The fact that our martian is deeply uncommunicative will not render it logically omniscient. We assumed that our martian calculates and reasons; it may have done the calculation that gave it the schmelief that seventeen is prime without yet having constructed the proof that would lead it to schmelieve that the sum of the angles of a Euclidean triangle is equal to a straight angle. If linguistic structure plays a role in the explanation of the fact that there can be equivalent but distinct objects of schmelief, we will have to explain that role in a way that does not involve the structure of outward expressions. Perhaps the same explanation will work for belief as well.

I have considered two ways in which semantic structure might be involved in belief. First, the informational content of a belief may be stored in a form that has a semantic structure. Second, beliefs may be expressed in a form that has semantic structure. These two ways in which semantic structure is involved are different: information might be stored in one form and expressed in a different one, or it might be stored in a linguistic form and not expressed at all. Even if one assumes that all beliefs are represented internally in linguistic form, it is not plausible, I have argued, to take belief attributions to be making a claim about this form. And while it might be plausible on some conceptions of belief to take belief attributions to be making a claim about the form in which beliefs are expressed, this won't provide a general solution to the problem of the ignorance of necessary equivalence. The real problem is that in case where an agent believes that P but fails to realize that Q where P and Q are necessarily equivalent, it seems that this failure is a fact about the information available to the agent, and

not just a fact about how he represents, either internally or externally, the information that he has. If this is right, then the question is, what information does an agent have, and what information does he lack, when he believes or presupposes one thing while not believing or presupposing something equivalent. This question is deflected, rather than answered, when we shift our attention to the means used to represent the information, and away from the information itself. In the remainder of this paper, I will consider two ways that semantic structure may be involved in the informational content of belief, and not just in the means of representing that content.

First, and most simply, beliefs may be *about* semantic structure, or about either internal or external representations that have semantic structure. One may have beliefs about the informational content of a representation, or about its truth value. One can believe that a certain statement has a certain content, or a certain semantic structure, without having an opinion about the truth value of the statement, and one can believe that a certain statement is true without having an opinion about what the statement says, or what information it conveys. Beliefs can be about sentences and speech acts in a public language, and they can be about internal representations, linguistic or otherwise. Some of our internal representations may be about other of our internal representations — about what they mean, and whether they are right.

How is semantic knowledge, ignorance and error relevant to the general problem of the object of belief? It is clear that some kind of semantic ignorance or error will be present in any prima facie case of failure to believe things with the same informational content. Suppose M and N are fine-grained objects of belief of some kind with the same informational content. Suppose x believes M, but not N. Then obviously, x will have to be ignorant of the fact that M and N have the same informational content, since he does not know that they are even materially equivalent. But then he must be ignorant, either of what content M has, or of what content N has, and this is purely semantic information.

So it seems clear that in the kind of case that motivates us to look for a more fine-granied object of belief — cases where it seems that an agent believes something without believing something necessarily equivalent to it — there is always a difference in the information

available to the agent, and not just a difference in the way the agent packages the information. Might it be that when we attribute belief that *P* and deny belief that *Q* in case where semantic theory tells us that *P* and *Q* are necessarily equivalent, the beliefs being attributed and denied are beliefs that are at least in part about semantic relations? Might the information that distinguishes between necessary truths (which on a straightforward interpretation all have the same informational content) be semantic information — information about the expressions and structures used to state those truths?

In some simple case, it should be uncontroversial that this is right: if O'Leary fails to believe that all woodchucks are groundhogs, or that a fortnight is a period of fourteen days, then it is clear that the information O'Leary lacks, and the information that we are saying he lacks when we deny that has those beliefs, is information about the semantic values of certain words. But in most cases, the information in question does not seem, intuitively, to be information about expressions. Plausible semantic theories tell us, for example, that it is necessarily true that Hesperus is Phosphorus, and that measles is caused by a virus. If I point at Oliver North and say "*That* is Oliver North," the proposition I express is necessarily true. And of course all mathematical statements are necessarily true or necessarily false. But none of these statements seems to be about language; the information they convey seems to be about astronomical and medical facts, facts about who is being pointed at, or, in the case of arithmetical and geometric statements, facts about numbers and the abstract structure of certain spaces. One does not need to know any names to know, or be ignorant of, the fact that Hesperus is Phosphorus. and we may share mathematical beliefs with those who do not share our language. The ancient Greeks, for example, believed that the square root of two was irrational, and so do we. How can these facts and intuitions be reconciled with the hypothesis that the information relevant to distinguishing necessarily equivalent beliefs, and beliefs in necessary truths, is information about semantic values and structures?

First, it may help to keep in mind that questions about semantic values and structures may be mixed with questions about the world. If I have certain partial information about the content of a statement, then learning that the statement is true may tell me something about the

world without giving me the information that the statement itself conveys. Suppose you say to me, "John is under the bed," and I know that you are referring either to John Jones or John Smith, but I don't know which. I accept that what you tell me is true, but I don't know what you have told me. I do, however, learn something about who is under the bed from what you told me. And I am now in a position to infer something that has nothing to do with language from a piece of purely semantic information: if I learn who you were referring to, I can infer from what I already know who it is that is under the bed.

Sometimes semantic values are determined as a function of certain extra-linguistic facts. According to causal theories of reference, who I am referring to with a proper name, and what I am saying when I use the name, may depend on facts about causal chains of which I am ignorant. Anyone who accepts such a theory of reference must agree that the fact that Hesperus is Phosphorus is a semantic fact in the following sense: one who learns that Hesperus is Phosphorus learns a fact on which the semantic value of the names "Hesperus" and "Phosphorus" depends; no one who is ignorant of the fact that Hesperus is Phosphorus can be fully informed about the informational content of statements containing the names "Hesperus" and "Phosphorus." But while this information is in this sense semantic, it is also astronomical. One who learns that Hesperus is Phosphorus learns something about the way the solar system is arranged.

Beliefs about semantics can be beliefs about specific expressions, speech acts, or representational tokens, but they can also be about more abstract structures shared by specific tokens. Suppose I know of a set of sentences — sentences that all have the same structure at some level of abstraction — that they have the same structure and semantic value. Suppose there is a different set of sentences that share a different structure but the same semantic value, as those in the first set. Suppose I also know of these sentences that they are equivalent to each other, but I do not know that these sentences are equivalent to those in the first set. We might describe this cognitive situation in terms of the structures shared by the sentences rather than in terms of the specific sentences that share the structure. Now suppose there is a different person who speaks a different language, but who is in a parallel cognitive situation. The sentences she has beliefs about will be different,

but the structures they share may be the same. There is a common piece of semantic information that these speakers of different languages lack: the information that representations with *this* structure have the same informational content as representations with *that* structure.[9] Even though the ancient Greeks shared no beliefs with us about the specific words, numerals and other notation that I would use to say that square root of two is irrational, it may be this kind of cognitive parallel that explains why we can correctly attribute this belief to them.

People who speak different languages and who use different notations may have the same mathematical beliefs when there are relevant similarities between the semantic and notational structures of the different languages, but it is sometimes difficult to separate the content of a mathematical belief from the means used to express it. Suppose there were a community of English speakers that grew up doing its arithmetic in a base eight notation. The words "eight" and "nine" don't exist in its dialect; the words "ten" and "eleven," like the numerals "10" and "11" denote the numbers eight and nine. Now suppose that some child in this community has a belief that he would express by saying "twenty-six times one hundred equals twenty-six hundred." Would it be correct to say that this child believes that twenty-two times sixty-four equals fourteen hundred eight? This does not seem to capture accurately his cognitive state. His belief, like our simple arithmetical beliefs, is not relly a belief about the numbers themselves, independently of how they are represented.

Mathematical information is most often received in linguistic form, and the behavior that mathematical beliefs dispose us to engage in is primarily behavior that involves linguistic and other representations: calculation, symbolic construction, and proof. It is not implausible, I think, to take representations and representational structures to be the subject matter of mathematics, and to be involved in the subject matter in other prima facie cases of necessarily equivalent but distinct objects of belief. If we do assume this, then we can distinguish between different mathematical truths and other equivalent statements and clauses, and we can explain how semantic structure is essential to some objects of belief, without giving up the idea that belief attributions relate a believer to a coarse-grained informational content.

Before concluding, I want to mention briefly a fourth way that

language may be involved in belief, a way that is independent both of the subject matter of the belief and of the way the believer represents or expresses his belief. Linguistic practices and institutions are essential to the medium through which a lot of information is transmitted. The point is not just that linguistic communication may be part of the means by which someone comes to be in a belief state; rather, the point is that facts about linguistic communication may be essential to the fact that certain belief states have the informational content that they have. Suppose I am told, and accept as true, something about a place I have never heard of before. To borrow an example discussed in several places by Daniel Dennett, suppose I am told that Balzac was married in Berdichev. Now one thing I acquire from this communication is the partly metalinguistic information the Balzac was married in a place called "Berdichev," but it seems plausible to say that I also come to believe something about Berdichev itself — that Balzac was married there. By the mere fact that I assent to and remember this sentence, I am in a state that tends, under normal conditions, to carry the information that Balzac was married in Berdichev. Under normal conditions, I would not assent to that statement unless it were true, since I would not have assented to the statement if someone hadn't made it, and under normal conditions what people say is true. We must assume that statements are normally true if we are to explain how information can be transmitted by linguistic communication.[10]

If this is right, then the representational state that constitutes my believing that Balzac was married in Berdichev will be a state that has *that* informational content only because of facts about the semantics of the public language. If "Berdichev," as used by my informant, had referred to Novokuznetsk then the same internal state would have been the belief that Balzac was married in Novokuznetsk.

This kind of case is different from each of the other three kinds of language dependence that I have discussed. First, while a belief that is language dependent in this way *might* be represented internally in a linguistic form — for example the English sentence itself might be the form that the representation took — this is not essential. Information received in linguistic form might instead be used to construct or modify a mental map. Second, people need not be disposed to express beliefs that are language dependent in this way: they may be unconscious or tacit. Third, the content of such beliefs need not involve language.

Even if my belief that Balzac was married in Berdichev is language dependent in this way, someone else — Balzac's wife, for example — might have believed the same thing without her belief having any dependence on linguistic facts at all.

Dennett, in one context where he discusses the Berdichev example, suggests that we should distinguish sharply a primitive, language independent kind of intentional state — belief — whose objects are informational contents, from a more sophisticated state — *assent* or *opinion* — whose objects are sentences "collected as true."[11] But I think we can better understand the phenomenon Dennett is pointing to, and get clearer about the way that assent and opinion interact with belief in general, if we see language dependent opinion as a special case of the general kind of belief that applies to nonspeaking animals as well as to ourselves.

Whatever else it is to have a belief, I have suggested, it is at least to be in a state that tends to carry information, where information is understood in terms of a pattern of counterfactual dependencies of internal states on states of that part of the world that determines the subject matter of the beliefs. For such information-carrying states to be belief states it is at least necessary that their informational content be accessible, that the state tend to dispose the believer to behave in ways that are systematically related to its content. One can characterize systems and states of this general kind independently of linguistic structure and linguistic practice, and one can explain many of the puzzling features of intentional states in terms of the fact that they are state of systems of this kind. But while there can be systems of this kind that don't involve linguistic representation, language is obviously involved in the case of the paradigms of believers — ourselves — and perhaps in the case of anything capable of being in intentional states that are rich and powerful enough to be called beliefs. My main point has been that language can play many different roles, and that the languages and linguistic representations involved at different points in such a system need not be the same. It helps to get clearer about the different roles that language can play in belief — the different ways that beliefs and belief attributions can be language dependent — to see believers as information carrying and information using systems of this kind. And if we can get clear about the different roles that language plays in the beliefs of sophisticated believers, I think we will see how to

reconcile the simple conception of content as informational content with the subtle distinctions that can be made between objects of beliefs.

NOTES

[1] In *Inquiry* (Cambridge, MA: Bradford books/The MIT Press, 1984).
[2] Davidson's paratactic account, and any quotational account of belief attribution is an instance of this strategy. See "On Saying That" in Donald Davidson, *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984), 93—108, and Christopher Hill, "Toward a Theory of Meaning for Belief Sentences," *Philosophical Studies, 30* (1976), 206—226.
[3] For an abstract and unconstrained theory of fine-grained objects of attitude, see R. H. Thomason, "A Model Theory for Propositional Attitudes," *Linguistics and Philosophy, 4* (1980), 47—70. For an account of objects of belief as sets of situations, see Jon Barwise and John Perry, *Situations and Attitudes* (Cambridge, MA: Bradford Books/ the MIT Press, 1983). For some criticisms of Barwise and Perry's account, see Scott Soames, "Lost Innocence," *Linguistics and Philosophy, 8* (1985), 59—72.
[4] For developments of this strategy, see Max Cresswell, *Structured Meanings* (Cambridge, MA: Bradford Books/the MIT Press, 1985), John Bigelow,"Believing in Semantics," *Linguistics and Philosophy, 2* (1978), 101—144, and Nathan Salmon, *Frege's Puzzle* (Cambrige, MA: Bradford Books/the MIT Press, 1986).
[5] See William Lycan, "Tacit Belief," in Radu Bogdan (ed.), *Belief: Form, Content, Function* (Oxford: Oxford University Press, 1986) and Robret Cummins, "Inexplicit Information," in Myles Brand and Robert Harnish (eds.), *The Representation of Knowledge* (Tucson: The University of Arizona Press, 1986), 116—126.
[6] Donald Davidson, "Belief and the Basis of Meaning," in *Inquiries into Truth and Interpretation*, p. 144.
[7] Donald Davidson, "Thought and Talk," in *Inquiries into Truth and Intepretation*, p. 170.
[8] Donald Davidson, "Belief and the Basis of Meaning," *Inquiries into Truth and Interpetation*, p. 144.
[9] There are some ambiguities one has to be careful about here. The relevant belief is not about the structure itself, but about certain sentences that in fact share the structure. The proposition that a certain structured meaning determines a certain informational content will be a necessary truth, and so we cannot use the informational content of this proposition to distinguish different mathematical beliefs.
[10] Of course being in a state that tends to carry the information that *P* is not sufficient for believing that *P*. Other conditions, including the condition that the information be accessible to guide action, need to be met. So we need not say that every case of accepting that a statement one hears or reads is true is a case of believing what that statement says.
[11] Daniel Dennett, "How to Change Your Mind," *Brainstorms* (Montgomery, VT: Bradford Books, 1978), chapter 16. Dennett is endorsing a suggestion made by Ronald de Sousa in "How to give a piece of your mind: or, the logic of belief and assent," *Review of Metaphysics, 25* (1971), 52—79.

Department of Linguistics and Philosophy,
Massachusetts Institute of Technology,
Cambridge, MA 02139,
U.S.A.