Notre Dame Journal of Formal Logic
Volume 23, Number 2, April 1982

Conceptual Role Semantics

GILBERT HARMAN*

In this paper, I will attempt to say something about a number of issues that arise in connection with "conceptual role semantics," the approach to semantics for which I have the most sympathy. On this occasion I will be able only to sketch aspects of this view without being able to give detailed arguments.

1 What is conceptual role semantics? A theory of conceptual role semantics involves the following two claims:

1. The meanings of linguistic expressions are determined by the contents of the concepts and thoughts they can be used to express.
2. The contents of concepts and thoughts are determined by their functional role in a person's psychology.

"Thoughts" here include beliefs, hopes, desires, fears, and other attitudes, in addition to thoughts properly so called. "Functional role" includes any special roles a concept may play in perception and in inference or reasoning, including practical reasoning that leads to action.

Conceptual role semantics represents one thing that might be meant by the slogan "meaning is use". But a proper appreciation of the point requires distinguishing (at least) two uses of symbols, their use in calculation, as in adding a column of figures, and their use in communication, as in telling someone the result.

1.1 Two uses of symbols: communication and calculation Symbols that are being used in calculation are typically not being used at that time for communication. When you add a column of figures you are not normally communicating anything even to yourself. A similar point holds in reverse.

*I am indebted to Scott Soames for his comments on this paper.

Received July 17, 1981

Normally you communicate the results of your calculation to someone else only after you are done calculating. There are, of course, mixed cases. You might go through a calculation on the blackboard, intending your audience to see how things come out.

Conceptual role semantics may be seen as a version of the theory that meaning is use, where the basic use of symbols is taken to be in calculation, not in communication, and where concepts are treated as symbols in a "language of thought". Clearly, the relevant use of such symbols—the use which determines their content—is their use in thought and calculation rather than in communication. If thought is like talking to yourself, it is the sort of talking involved in figuring something out, not the sort of talking involved in communication. Thinking is not communicating with yourself.

However, it would be more accurate to say content is use than to say meaning is use; strictly speaking, thoughts and concepts have content, not meaning.

1.2 The meaning of "meaning" I assume, following Grice [6], that we can distinguish what he calls natural meaning (smoke means fire) from what he calls nonnatural meaning (the German word "Feuer" means fire), and we can also distinguish (nonnatural) speaker or user meaning (what a speaker or user of certain symbols means) from what certain words, expressions, or other symbols mean.

Grice proposes to analyze expression meaning in terms of speaker meaning; and he proposes, more controversially, to analyze speaker meaning in terms of a speaker's intentions to communicate something. This last proposal appears to overlook the meaningful use of symbols in calculation. You might invent a special notation in order to work out a certain sort of problem. It would be quite proper to say that by a given symbol you meant so-and-so, even though you have no intentions to use these symbols in any sort of communication.

There does seem to be some sort of connection between speaker or user meaning and the speaker's or user's intentions. Suppose you use your special notation to work out a specific problem. You formulate the assumptions of the problem in your notation, do some calculating, and end up with a meaningful result in that notation. It would be correct to say of you that, when you wrote down a particular assumption in your notation, you meant such and such by what you wrote: but it would be incorrect to say of you that, when you wrote the conclusion you reached in your notation, you meant so and so by what you wrote. This seems connected with the fact that, in formulating the assumption as you did in your notation, you intended to express such and such an assumption; whereas, in reaching the conclusion you reached in your notation, your intention was not to express such and such a conclusion but rather to reach whatever conclusion in your notation followed from earlier steps by the rules of your calculus. This suggests that you meant that so and so in using certain symbols if and only if you used those symbols to express the thought that so and so, with the intention of expressing such a thought.

Unexpressed thoughts (beliefs, fears, desires, etc.) do not involve meaning. We would not ordinarily say that in thinking as you did you meant that so and so. If thoughts are in a language of thought, they are not (normally) expressed

in that language. (I say "normally" to allow for the possibility of thoughts that are thought in a language used in communication, thoughts in English for example.)

Concepts and other aspects of mental representation have content but not (normally) meaning (unless they are also expressions in a language used in communication). We would not normally say that your concept of redness meant anything in the way that the word "red" in English means something. Nor would we say that you meant anything by that concept on a particular occasion of its exercise.

2 Content and inferential role

2.1 Types of role Assuming conceptual role semantics as a basic framework, it is plausible that all concepts have a function in reasoning that is relevant to their content. No doubt, some concepts have the content they have primarily because of a special role they play in perception—color concepts for example. But the content of even these concepts depends to some extent on inferential role. A given color concept is the concept of a normally persisting characteristic of objects in the world, a characteristic that can be used both to keep track of objects and as a sign of other things. For example, greenness is a sign of unripeness in certain fruits. Moreover, there are various internal relations among colors. From the premise that an object is red at a certain place at a certain time one can infer that the object is not also green at that place and time.

In the case of concepts of shape and number, inferential connections play a larger role. Perceptual connections are still relevant; to some extent your concept of a triangle involves your notion of what a triangle looks like and your concept of various natural numbers is connected with your ability to count objects you perceive. But the role these notions play in inference looms larger.

The concept expressed by the word "because" plays an important role in one's understanding of phenomena and has (I claim) a central role in inference, since inference is often inference to the best explanation. This role makes the concept expressed by "because" the concept it is, I believe. Is perception relevant at all here? Perhaps. It may be that you sometimes directly perceive causality or certain explanatory relations, and it may be that this helps to determine the content of the concept you express with the word "because". Or perhaps not. Maybe the perception of causality and other explanatory relations is always mediated by inference.

Logical words like "and", "not", "every", and "some" express concepts whose function in inference seems clearly quite important to their content, which is why it seems so plausible to say that these words do not mean in intuitionistic logic and quantum logic what they mean in so-called classical logic, although even here there may be crucial perceptual functions. It may, for example, be central to your concept of negation that you can sometimes perceive that certain things are not so, as when you perceive that Pierre is not in the cafe. It may be central to your concept of generality or universal quantification that you can sometimes perceive that everything of a certain sort is so and so, for instance that everyone in the room is wearing a hat.

It is possible that there are certain sorts of theoretical terms, like "quark", that play no role in perception at all, so that the content of the concepts they express is determined entirely by inferential role. (Even here maybe it is important to the concept of a quark that the concept should play a role in the perception of certain pictures or diagrams!)

2.2 Inference and implication Logical words have a function in inference and reasoning because certain implications depend on them. Inference is, of course, a process of thought which typically culminates in a change in view, a change in your beliefs if it is theoretical reasoning, a change in plans and intentions for practical reasoning. (There is also the limiting case in which you make no change.)

There is as yet no substantial theory of inference or reasoning. To be sure, logic is well developed; but logic is not a theory of inference or reasoning. Logic is a theory of implication or of argument, where an argument is conceived as a sequence of implications.

Logic is relevant to reasoning because implication is. Implication is relevant to reasoning because implication is an explanatory relation and, in reasoning, you try among other things to increase the explanatory coherence of your view (see [13]).

2.3 Logical form Accounts of the logical forms of sentences of a natural language can shed light on meaning to the extent that they indicate aspects of language on which implications may depend, since this is to indicate something about the inferential role played by the concepts expressed by those aspects of language.

Presumably, such accounts of logical form should be relevant to or perhaps even part of a grammatical analysis of the relevant sentences. Here is an area where there may be useful interaction between what philosophers do and what linguists do. However, as Chomsky [3] observes, distinctions that are important for linguistics may not coincide with the distinctions that are important for philosophers. Or, to put the point in another way, the factors that determine relations of implication may not all be of the same sort. Some may be aspects of what Chomsky calls "sentence grammar", others may not. And some aspects of "sentence grammar" that function syntactically like logical features may not be directly connected with implication. For example, Chomsky suggests that the rules of grammar that determine how quantifiers are understood, which are of course crucial in determining what the logical implications of the sentence are, may be the same as the rules that determine such things as the "focus" of a sentence, something which seems not to affect the logical implications of a sentence but only its "conversational implicatures".

2.4 Indeterminacy There are apparently competing analyses of the logical forms of sentences. Where one analysis sees modal logic or tense logic, another sees reference to possible worlds or times. Where one analysis sees reference to events, another analysis invokes an adverbial logic. And so on. Similarly, there are apparently competing grammatical theories: Montague grammar, Chomsky's current framework in terms of rules of government and binding, and many other variations. What are we to make of this?

Quine [22] argues plausibly that even all possible observations about a language may not decide between various locally incompatible "analytical hypotheses" (where by "analytical hypotheses" he means hypotheses about logical or grammatical form of the sort just mentioned). There has been considerable dispute as to exactly how Quine's claim should be interpreted, whether it is true, and what the implications of its truth might be. It has been said (falsely) that all Quine's thesis amounts to is the claim that a theory is underdetermined by the evidence. It has also been said (correctly, I believe) that whatever valid point Quine may be making, it does not involve any significant difference between the "hard sciences", like physics, and the study of language.

One issue suggested by Quine's argument is this. Suppose you have a theory, of physical reality or of language, which you think is true. Even though you think the theory is true, you can go on to consider what aspects of the theory correspond to reality and what aspects are instead mere artifacts of the notation in which the theory is presented. A true geographical description of the Earth will mention longitudes as well as cities and mountains, but longitudes do not have geographical reality in the way that cities and mountains do. It is true that Greenwich, England, is at zero degrees longitude, but this truth is an artifact of our way of describing the Earth, since there are other equally true ways of describing the geography of the Earth that would assign Greenwich other longitudes. Similarly, there are various true physical descriptions of the world, which assign a given space-time point different coordinates. It may be true that under a particular description a particular point has the special coordinates (0,0,0,0) but that is an artifact of the description which, by itself, does not correspond to anything in reality. And the same is true of grammars and theories of logical form. Even if a given account of grammar or logical form is true, there is still a question of which aspects of the account correspond to reality and which aspects are merely artifacts of that particular description. It is quite possible that several different locally incompatible accounts might all be true, just as several different locally incompatible assignments of longitudes and latitudes to places on Earth might all be true.

This might be put in another way. Reality is what is invariant among true theories. Geographical reality is what is invariant in different true geographical descriptions of the world. Physical reality is what is invariant in different true physical descriptions of the universe. What worries Quine is that he has a fairly good sense of physical and geographical reality but little or no sense of grammatical reality or of the reality described by accounts of logical form. Indeed, Quine is inclined to think that there are only two possibilities here:

1. physical reality at the level of neurophysiology
2. behavioral reality, including dispositions to behave in various ways.

An alternative view is that there are other, functionally defined levels of reality between the two levels Quine acknowledges. I see no other way to investigate this issue except by seeing where current investigation of grammar and logical form ultimately leads. Of course, from a heuristic point of view, it is probably best to suppose that different accounts of grammar and logical form make conflicting claims about reality unless there is some reason to think otherwise (see [16]).

2.5 Meaning and truth conditions Davidson [4], Lewis [17], and others have argued that an account of the truth conditions of sentences of a language can serve as an account of the meanings of those sentences. But this seems wrong. Of course, if you know the meaning in your language of the sentence *S*, and you know what the word "true" means, then you will also know something of the form "*S* is true if and only if . . ."; for example "'Snow is white' is true if and only if snow is white" or "'I am sick' is true if and only if the speaker is sick at the time of of utterance". But this is a trivial point about the meaning of "true", not a deep point about meaning (see [10]).

This is not to deny that attempts to develop theories of truth adequate for certain aspects of natural language may well shed light on meaning. Examples might include the truth functional analysis of "and", "not", and "or"; the Frege-Tarski analysis of quantification; Davidson's analysis of action sentences; and the possible worlds account of modality. But in all these cases the analyses help specify implications among sentences. Their bearing on meaning may be due entirely to that, apart from anything further necessary to having a theory of truth—although this, of course, allows that there may also be a heuristic point to attempting to develop theories of truth (see [9] and [16]).

2.6 Probabilistic semantics Field [5] suggests that inferential role might be captured in terms of a probability distribution. This would yield at best a theory of probabilistic implication or coherence, not a theory of inference in the relevant sense, involving (normally) a certain sort of change in view. Furthermore, people do not and could not operate probabilistically, since keeping track of probabilities involves memory and calculating capacities which are exponentially exploding functions of the number of logically unrelated propositions involved (see Harman [11] and [14]).

For the most part you have to accept propositions in an all or nothing way. Conservatism is important. You should continue to believe as you do in the absence of any special reason to doubt your views, and in reasoning you should try to minimize change in your initial opinions in attaining other goals of reasoning. Such other goals include explanatory coherence and, of course, practical success in satisfying your needs and desires (see [11] and [12]). But these points are vague and do not take us very far. Furthermore, something ultimately needs to be said about practical reasoning (see [15]).

3 Conceptual role and external world Conceptual role semantics does not involve a "solipsistic" theory of the content of thoughts. There is no suggestion that content depends only on functional relations among thoughts and concepts, such as the role a particular concept plays in inference. (Field [5] misses this point.) Also relevant are functional relations to the external world in connection with perception, on the one hand, and action, on the other. What makes something the concept red is in part the way in which the concept is involved in the perception of red objects in the external world. What makes something the concept of danger is in part the way in which the concept is involved in thoughts that affect action in certain ways.

3.1 The division of linguistic labor The content of certain concepts appears to depend crucially on functional relations between those concepts and certain

words in a public language. You might have a concept of an oak tree by virtue of which you have thoughts about oak trees where the crucial functional relation is a relation between your concept and the word "oak" in English. You might for example wonder whether there were any oak trees in your back yard even if you cannot distinguish oak trees from elm trees and do not know any of the distinguishing properties of these two sorts of trees (see Putnam [21]).

Conceptual role semantics asserts that an account of the content of thoughts is more basic than an account of communicated meaning and the significance of speech acts. In this view, the content of linguistic expressions derives from the contents of thoughts they can be used to express. But allowance must also be made for cases in which the content of your thoughts depends in part on the content of certain words, such as "oak" and "elm".

Of course, in this case, there are other people who can recognize oaks and distinguish them from elms and who know various distinguishing properties of the trees. These other people may have a concept of an oak tree which has functional roles that are sufficient to make it the concept of an oak tree apart from any relations the concept has with the word "oak". It is plausible

1. that their concept acquires its content from this aspect of its functional role, i.e., its role apart from its relation to the word "oak"
2. that the word "oak" as they use it has the content it has because of its connection with their concept of an oak tree
3. that the word "oak" as used by a more ignorant person can have the same content by virtue of connections between that person's ignorant use of the word and the expert's use
4. that the content of the more ignorant person's concept of an oak tree derives from its connection to his or her use of the word.

This would still allow one to say that the meanings of words derive ultimately from the contents of concepts the words are used to express, where the contents of these concepts do not themselves derive from the meanings of words; however the meanings of a particular person's words may not derive in this way from the contents of that person's concepts.

This suggests an interesting question. Is there any word for which there is a real division of linguistic labor, so that no single person has a corresponding concept whose content is functionally determined apart from its relation to the person's use of that word? It is certainly imaginable that this should be so in connection with some sort of group investigation. Different people might investigate different aspects of a phenomenon which each might identify as "whatever it is we are all investigating and which has such and such effects when investigated in the way in which I have investigated it". Even in such a case, the meaning of the word would derive from the role the corresponding concept plays in thought, although different aspects of that role would be fulfilled by different people's instance of the concept.

3.2 Content as relative to choice of a normal context Putnam imagines a world, which he calls "Twin Earth", which is just like Earth except for certain minor differences. There are on Twin Earth duplicates of all the people on Earth and the people on Twin Earth speak the same languages as on Earth,

using expressions in the same way, except that, because of the minor differences between Earth and Twin Earth, they sometimes refer to different things by their words. In particular, the main difference between Twin Earth and Earth is that where there is water on Earth there is on Twin Earth a liquid with the same macroproperties as water but a different chemical structure, which Putnam calls "XYZ".

Now, comparing Earth in 1750 (before the microstructure of water has been investigated) with Twin Earth at the corresponding time, we find that the English word "water" means something different in the two places, simply because the word is used on Earth to refer to what is in fact H_2O and is used on Twin Earth to refer to what is in fact XYZ. Similarly, where Earthlings think about H_2O, Twin Earthlings think about XYZ. This difference is not in 1750 reflected in any difference in dispositions to react to various perceptual situations, in any difference in inferences that people in the respective places would make, nor in any differences in the actions which people undertake as the result of thoughts involving the relevant concept.

The difference is also not simply a difference in context of utterance or context of thought. Suppose an Earthling were to travel by spaceship and land on an ocean of XYZ in Twin Earth. Looking around, the Earthling comes to believe there is water all around. This belief is false, since the Earthling's concept of water is a concept of something that is in fact H_2O. The Earthling's concept of water remains a concept of the same thing referred to by "water" on Earth even though the Earthling is now located in a different context. The context of the thoughts of the Earthling and the context of the thoughts of the Twin Earthlings are now the same; but their thoughts are about XYZ whereas his are still about water. So this difference in the content of the thoughts of Earthlings and Twin Earthlings cannot be simply a difference in the context in which they have their thoughts.

The difference is due rather, I suggest, to the fact that the content of a person's concept is determined by its functional role in some normal context. The normal context for an Earthling's thoughts about what he or she calls "water" is here on Earth, and the normal context for a Twin Earthling's thoughts about what he or she calls "water" is on Twin Earth.

The normal context can change. If the traveler from Earth to Twin Earth stays on, after a while the normal context for the concepts he or she uses will be properly taken to be the Twin Earthean context. Thoughts about what he or she calls "water" will be properly considered thoughts about XYZ rather than H_2O. There is of course a certain amount of arbitrariness in any decision about when this change has occurred. It will sometimes be possible with equal justice to consider a given thought to be a thought about H_2O or a thought about XYZ.

A similar arbitrariness would arise concerning a person created spontaneously in outer space as the improbable result of random events at the quantum level, supposing the person were saved from space death by a fortuitously passing space ship, and supposing the person spoke something that sounded very much like English. Suppose, indeed, that this person is a duplicate of you and also (of course) of your Twin Earth counterpart. When the person has thoughts that he or she would express using the term "water", are these thoughts about water (H_2O) or thoughts about XYZ? If we interpret this person's thoughts against a normal background on Earth, we will interpret the

relevant thoughts as thoughts about water. If we take the normal background to be Twin Earth, they are thoughts about XYZ. Clearly it is quite arbitrary what we say here.

3.3 Qualia According to conceptual role semantics, then, the content of a thought is not a matter of the "intrinsic nature" of either that thought or other mental states and experiences but is rather a matter of how mental states are related to each other, to things in the external world, and to things in a context understood as a normal context. There is a familiar objection to this (Block and Fodor [2] and Nagel [19]) which claims that content is not determined always by such functions or relations. In this view the intrinsic qualities or "qualia" of certain experiences are sometimes relevant. It is said that your concept of red involves your notion of what it is like to see something red, where what it is like to see something red is not just a matter of the functional or relational characteristics of the relevant experience but of its intrinsic character as well.

One argument for this is that it is possible to imagine a person whose spectrum was inverted with respect to yours, so that the quality of experience you have in seeing something red is the quality this other person has in seeing something green, the quality of experience you have in seeing something blue is the quality this other person has in seeing something orange, and similarly for other colors, although in all relevant respects your color experiences function similarly, so that each of you is just as good as the other in applying the public color words to colored objects. According to this argument, the two of you would have different concepts which you would express using the word "red", although it might be difficult or even impossible to discover this difference, since it is not a functional difference.

I speak of an "argument" here, although (as D. Lewis [18] has observed in a similar context) the "argument" really comes down simply to denying the functionalist account of the content of concepts and thoughts, without actually offering any reason for that denial. This makes the argument difficult to answer. All one can do is look more closely at a functionalist account of the content of color concepts in order to bring out the way in which, according to functionalism, this content does not depend on the intrinsic character of experiences of color.

How could you imagine someone whose spectrum was inverted with respect to yours? One way would be to imagine this happening to yourself. Suppose there were color-inverting contact lenses. You put on a pair of lenses and the colors things seem to have are reversed. The sky now looks orange rather than blue, ripe apples look green, unripe apples look red, and so on. Suppose you keep these lenses on and adapt your behavior. You learn to say "green" rather than "red" when you see something that looks the way red things used to look; you learn to treat what you used to consider a green appearance of apples as a sign of ripeness, and so on. The years pass and your adaption becomes habitual. Wouldn't this be an intelligible case in which someone, the imagined future you, has a notion of what it is like to have the experience of seeing something to which the term "red" applies, where the notion functions in exactly the way in which your notion of what such an experience is like functions, although your notions are different? The functionalist must deny this and say that the imagined you associates the same

concept with the word "red" as the actual you does now and indeed sees the world as you now do.

Consider an analogous case. There actually exist lenses that are spatially inverting. With these lenses on, things that are up look down and vice versa. At first it is very difficult to get around if you are wearing such lenses, since things are not where they seem to be. But after a while you begin to adapt. If you want to grab something that looks high, you reach low, and vice versa. If you want to look directly at something that appears in the bottom of your visual field you look up. And so on. Eventually, such adaption becomes more or less habitual.

Now functionalism implies that if you become perfectly adapted to such space inverting lenses, then your experience will be the same as that of someone who is not wearing the inverting lenses (who has adapted to not wearing them if necessary), because now the normal context in relation to which your concepts function will have become a context in which you are wearing the inverting lenses. And in fact, people who have worn such lenses do say that, as they adapt to the lenses, the world tends to look right side up again (see [23], [20], and [24]).

Similarly, functionalism implies that if you become perfectly adapted to color-inverting lenses, the world will come to look to you as it looked before in the sense that given such perfect adaption the normal context in which your color concepts function will be a context in which you are wearing the color-inverting lenses. According to functionalism, the way things look to you is a relational characteristic of your experience, not part of its intrinsic character.

In order to get a feel for this aspect of conceptual role semantics, it may be useful to consider certain further cases. Consider Inverted Earth, a world just like ours, with duplicates of us, with the sole difference that there the actual colors of objects are the opposite of what they are here. The sky is orange, ripe apples are green, etc. The inhabitants of Inverted Earth speak something that sounds like English, except that they say the sky is "blue", they call ripe apples "red", and so on. Question: what color does their sky look to them? Answer: it looks orange. The concept they express with the word "blue" plays a relevantly special role in the normal perception of things that are actually orange.

Suppose there is a distinctive physical basis for each different color experience. Suppose also that the physical basis for the experience of red is the same for all normal people not adapted to color inverting lenses, and similarly for the other colors. According to conceptual role semantics this fact is irrelevant. The person who has perfectly adapted to color inverting lenses will be different from everyone else as regards the physical basis of his or her experience of red, but that will not affect the quality of his or her experience.

Consider someone on Inverted Earth who perfectly adapts to color inverting lenses. Looking at the sky of Inverted Earth, this person has an experience of color whose physical basis is the same as that of a normal person on Earth looking at Earth's sky. But the sky looks orange to the person on Inverted Earth and blue to normal people on Earth. What makes an experience the experience of something looking the color it looks is not its intrinsic character and/or physical basis but rather its functional characteristics within an assumed normal context.

Consider a brain spontaneously created in space as the improbable result of random events at the quantum level. The physical events in the brain happen to be the same as those in you on Earth looking at the sky on a clear day and also the same as those in a person adapted to color inverting spectacles looking at the sky of Inverted Earth. What is it like for the brain? Is it having an experience of orange or of blue? According to conceptual role semantics, there is no nonarbitrary way to answer this question; it depends on what is taken as the normal context for assessing the functional role of events in that brain. If the normal context is taken to be the normal context for perception of color on Earth, the brain is having an experience of blue. If the normal context is taken to be the normal context for a wearer of inverted spectacles on Inverted Earth, the brain is having an experience of orange.

4 Meaning and speech acts It may be an exaggeration to say that all aspects of linguistic meaning derive from the use of language to express concepts with corresponding contents. Certain aspects of meaning seem to derive directly from the use of language in speech-acts and communication. Conceptual role semantics may therefore have to allow for this, attempting (of course) to treat it as a minor and relatively peripheral phenomenon. On the other hand, it may be that no such concession needs to be made. Here are some tentative and preliminary reflections on some of the issues involved.

4.1 Greetings The cases that may seem clearest here concern expressions which are used primarily in greetings and salutations, words like "hello" and "good-bye". The meanings of these words do not seem to derive from any use they might have to express single concepts that play a distinctive role in calculation and thought. It is true that you might, on occasion, "greet" a new idea or percept with the thought, "Hello! What's this?" But this seems to be a case in which the content of your thought derives from the use of the word "Hello" to greet people and begin conversations rather than the other way around.

However, it may be that such words can be analyzed as expressing a combination of concepts which individually have contents connected with distinctive conceptual roles. For example, perhaps "hello" means something like "I acknowledge your presence", or sometimes maybe "let us talk", and analogously for "good-bye" and other words and phrases of this sort. If so, the issue becomes whether the aspect of meaning expressed by the imperative in "let us talk" and the performative aspect of the meaning of "I acknowledge your presence" derive irreducibly from the use of words in speech acts or ultimately derive instead from the use of language to express concepts whose content is determined by their role in calculation and thought. That is a complex issue which we must consider in a moment. As for the question of whether this is the right way to analyze "hello", "good-bye", and so on, I am not very sure what to think. The analyses I have suggested seem to leave something out, but this might be accommodated by better analyses.

4.2 Words of politeness Before turning to imperatives and performatives, we might consider words and phrases that function as forms of deference and politeness: "please" and "thank you", for example. The use of such words and phrases seems to presuppose some sort of social interaction, rather than simply expressing concepts with a content determined by the way these concepts

function in thought and calculation. It is true that you might say "please" or "thank you" to yourself, in thought, if you were speaking to yourself as an instructor or coach might. But this seems a very special case, which itself presupposes the social use of these expressions.

Again, it might be argued that in this use "please" means the same as some longer expression, such as "if you please" or "if it pleases you to do so", and "thank you" means something like "I hereby thank you". If so, the issue would become whether the meaning of the term "you" derives entirely from its use to express concepts whose content is determined by their distinctive function in calculation or thought, an issue we need to discuss. Also, there is again the question of whether the performative element, in "I hereby thank you", carries a meaning which at least in part is irreducibly connected with speech acts, another issue we will be coming to in a moment. As for the question whether such polite phrases can always be analyzed in this sort of way, I am not sure what to say.

4.3 "You" What about the meaning of the word "you"? It is plausible that this word means something like "the person I am now addressing" and each of the words in this phrase is plausibly held to have a meaning that depends on the concept the word is used to express, where the concept in question has the content it has by virtue of the way it functions in thought and calculation. So, even though the word "you" has the distinctive function in speech acts and communication of designating the intended audience, this function can plausibly be explained in terms of the functional roles in thought of the concepts the word is used to express.

4.4 Imperative and interrogative mood The use of the imperative mood in English seems to carry a certain meaning, connected with the giving of directions of some sort. This does not by itself imply this meaning does not derive from the content of certain concepts with a distinctive use in thought and calculation. Indeed, something in your thoughts functions to distinguish your beliefs from your plans and intentions, which are directions of a sort. So, there is a sense in which the upshot of practical reasoning is a modification of certain directions one intends to follow. Perhaps the imperative mood serves to express the concept which functions in thought to thus distinguish practical or directing thoughts from theoretical thoughts.

Similar remarks apply to the interrogative mood. Indeed, questions are not unlike requests for information, so that the interrogative mood is plausibly analyzed in terms of the imperative mood. In any event, questions obviously have a function in thinking. You pose a problem to yourself and work out the answer, perhaps by posing various subquestions and answering them.

4.5 Performatives Consider next explicit performatives, like "I promise to be there" and "I hereby apologize for my rude behavior". It seems part of the meaning of such sentences that they are used not to describe the speaker as promising or apologizing but actually to do the promising or apologizing. Furthermore, it is plausible that promising or apologizing to yourself is not a typical or normal case of promising or apologizing and is rather the sort of case which is to be understood in relation to more typical or normal cases in which you promise someone else something or apologize to another person.

On the other hand, each of the words in a sentence like "I promise to be there" has a meaning which expresses a concept whose content is arguably determined by its functional role in calculation and thought. And it is possible that the meaning of the whole sentence, including whatever gives the sentence its performative function of being appropriate for actually promising, arises from the meaning of the words used in a regular way. Given what the words in the sentence mean and given the way these words are put together, it may be predictable that the sentence has a performative use (see [1]).

Suppose we adopted the convention that promises have to be made in some special way, for example by writing down the content of the promise in purple chalk on a special promise board that is not used for any other purpose. The convention would be that nothing else is to count as a promise. In such a case, the words, "I promise to be there" could not be used to promise you will be there. Would this be a way to pry off the performative meaning of "I promise" from that aspect of its meaning that derives from its use to express concepts whose content is determined by their functional role in thought? Not obviously. For one thing, this might change the concept of promising in a significant way. The word "promise" might not mean what it meant when a promise could be made by saying "I promise". It could be argued that if "I promise" means what it ordinarily means, then it follows from the concepts expressed by the words "I promise" that these words can be used to promise.

Alternatively, it might be said that, even if the word "promise" would retain (enough of) its usual meaning when promising was restricted by such a convention, the example is like one in which a special convention is adopted that an utterance of the sentence "The sky is blue" is not to be interpreted as an assertion that the sky is blue but rather as a question asking whether it rained last week. This would not show that there is any aspect of the meaning of an ordinary assertion of the sentence "The sky is blue", as we use it now without such a bizarre convention, which does not derive from the way the words in the sentence are used to express concepts that have the content they have because of their functional role in thought.

4.6 Conversational implicature Grice ([7] and [8]) argues that the implications of an utterance do not always correspond directly to the meanings of the linguistic expressions used, even in quite ordinary cases. He suggests in particular that what seem to be aspects of meaning may be due to "conversational implicature," i.e., to conclusions the audience is intended to reach by reflecting on the speaker's reasons for saying what is said, assuming the speaker is trying to be helpful. For example, if you use "either . . . or", as in "Albert is in either Boston or New York", you normally imply you do not know which. This does not have to be taken as showing something special about the meaning of "either . . . or", a difference in the meaning of this ordinary expression as compared with what the logician takes it to mean. Instead we can suppose this implication is due to the natural assumption that normally, if you know which city Albert is in, you will say which it is. To take another example, the apparent difference in meaning between "Mary closed the door and turned on the light" and "Mary turned on the light and closed the door" does not have to be explained by supposing that the English word "and" sometimes means "and then". The suggestion is that the difference can be explained by supposing that

a helpful speaker will normally relate events in an orderly way, so that a hearer is normally justified in supposing that the order in which the speaker relates the events is to be understood as the order in which they occurred. Similarly, Grice suggests that certain aspects of presupposition might be explained by considering the normal expectations of speakers and hearers.

However, some of these phenomena can occur in thinking to oneself, where they are presumably not due to conversational implicature. Calculation and reasoning often involve various presuppositions. One will normally want descriptions used in reasoning to relate events in an orderly way, so the same phenomenon with the word "and" may occur. On the other hand, it is doubtful that use of an "either . . . or" proposition in thought normally carries the implication that one does not know which alternative is the case; so *this* phenomenon may really occur only at the level of conversation.

*4.7 **Figurative language*** Metaphor and simile occur in thinking. I am inclined to think irony does not. I am not sure about hyperbole.

*5 **Conclusion*** To summarize: There are two uses of symbols, in communication and speech acts and in calculation and thought. Conceptual role semantics takes the second use to be the basic one. The ultimate source of meaning or content is the functional role symbols play in thought.

The content of a concept depends on its role in inference and sometimes in perception. Particularly important are a term's implications. Implication is relevant to inference and therefore to meaning, because implication is explanatory and inference aims at explanatory coherence. Accounts of truth conditions can shed light on meaning to the extent that they bring out implications; it is doubtful whether such accounts have any further bearing on meaning, although they may have heuristic value for studies of logical form. Probabilistic semantics does not provide an adequate conceptual role semantics, because people do not and cannot make much use of probabilitistic reasoning.

Allowance must be made for various connections between concepts and the external world. Some concepts have the content they have because of the words they are associated with, although (according to conceptual role semantics) this content ultimately always derives from someone's use on concepts. The content of concepts is often relative to a choice of a normal context of functioning. This is true of color concepts, despite the unargued view of some philosophers that these concepts depend on the intrinsic character of experience.

Finally, it is not clear whether any aspects of meaning derive directly from the use of language in speech acts in a way not reducible to the expression of concepts whose content is independently determined. In any event, many phenomena often taken to be particularly connected with speech acts and conversation also occur in calculation and thought.

REFERENCES

[1] Bach, K., "Performatives are statements, too," *Philosophical Studies,* vol. 28 (1975), pp. 229-236.

[2] Block, N. and J. A. Fodor, "What psychological states are not," *Philosophical Review,* vol. 81 (1972), pp. 159-181.

[3] Chomsky, N., *Reflections on Language,* Columbia University Press, New York, 1980.

[4] Davidson, D., "Truth and meaning," *Synthese,* vol. 17 (1967), pp. 304-323.

[5] Field, H., "Logic, meaning, and conceptual role," *Journal of Philosophy,* vol. 74 (1977), pp. 379-409.

[6] Grice, H. P., "Meaning," *Philosophical Review,* vol. 68 (1959), pp. 377-388.

[7] Grice, H. P., "The causal theory of perception," *Proceedings of the Aristotelian Society,* Supplementary Volume 35 (1961).

[8] Grice, H. P., "Logic and conversation," in *The Logic of Grammar,* eds., D. Davidson and G. Harman, Dickenson Press, Encino, California, 1975.

[9] Harman, G., "Logical form," *Foundations of Language,* vol. 9 (1972), pp. 38-65.

[10] Harman, G., "Meaning and semantics," in *Semantics and Philosophy,* eds., M. K. Munitz and P. Unger, New York University Press, New York, 1974.

[11] Harman, G., "Reasoning and explanatory coherence," *American Philosophical Quarterly,* vol. 17 (1980), pp. 151-157.

[12] Harman, G., "Reasoning and evidence one does not possess," *Midwest Studies in Philosophy,* vol. 5 (1980), pp. 163-182.

[13] Harman, G., "Logic and reasoning," to appear in the proceedings of a conference on logic held in New York in 1980.

[14] Harman, G., "Problems with probabilistic semantics," to appear in *Developments in Semantics,* eds., A. Orenstein et al., Haven Press, New York.

[15] Harman, G., "Willing and intending," to appear in a Festshrift for H. P. Grice, edited by R. Grandy and R. Warner.

[16] Harman, G., "Quine's grammar," to appear in *The Philosphy of W. V. Quine,* ed., P. Schillp, Open Court Publishing Company, La Salle, Illinois.

[17] Lewis, D., "General semantics," in *Semantics of Natural Language,* eds., D. Davidson and G. Harman, D. Reidel Publishing Company, Dordrecht, Holland, 1972.

[18] Lewis, D., "Mad pain and Martian pain," in *Readings in Philosophy of Psychology,* ed., N. Block, Harvard University Press, Cambridge, Massachusetts, 1980.

[19] Nagel, T., "What is it like to be a bat?," *Philosophical Review,* vol. 83 (1974), pp. 435-450.

[20] Pitcher, G., *A Theory of Perception,* Princeton University Press, Princeton, New Jersey, 1971.

[21] Putnam, H., "The meaning of meaning," in *Mind, Language, and Reality: Philosophical Papers,* Vol. 2, Cambridge University Press, Cambridge, Massachusetts, 1975.

[22] Quine, W. V., *Word and Object,* Massachusetts Institute of Technology Press, Cambridge, Massachusetts, 1960.

[23] Taylor, J. G., *The Behavioral Basis of Perception,* Yale University Press, New Haven, Connecticut, 1962.

[24] Thomas, S., *The Formal Mechanics of Mind,* Cornell University Press, Ithaca, New York, 1978.

Department of Philosophy
Princeton University
Princeton, New Jersey 08544