

LITERAL MEANING¹

I

Most philosophers and linguists accept a certain conception of the notion of the literal meaning of words and sentences and the relation between literal meaning and other semantic notions such as ambiguity, metaphor, and truth. In this paper I want to challenge one aspect of this received opinion, the view that for every sentence the literal meaning of the sentence can be construed as the meaning it has independently of any context whatever. I shall argue that in general the notion of the literal meaning of a sentence only has application relative to a set of contextual or background assumptions and finally I shall examine some of the implications of this alternative view. The view I shall be attacking is sometimes expressed by saying that the literal meaning of a sentence is the meaning that it has in the 'zero context' or the 'null context'. I shall argue that for a large class of sentences there is no such thing as the zero or null context for the interpretation of sentences, and that as far as our semantic competence is concerned we understand the meaning of such sentences only against a set of background assumptions about the contexts in which the sentence could be appropriately uttered.

I begin by stating what I take to be the received opinion as a set of propositions:

Sentences have literal meanings. The literal meaning of a sentence is entirely determined by the meanings of its component words (or morphemes) and the syntactical rules according to which these elements are combined. A sentence may have more than one literal meaning (ambiguity) or its literal meaning may be defective or uninterpretable (nonsense).

The literal meaning of a sentence needs to be sharply distinguished from what a speaker means by the sentence when he utters it to perform a speech act, for the speaker's utterance meaning may depart from the literal sentence meaning in a variety of ways. For example, in uttering a sentence a speaker may mean something different from what the

sentence means, as in the case of metaphor; or he may even mean the opposite of what the sentence means, as in the case of irony; or he may mean what the sentence means but mean something more as well, as in the case of conversational implications and indirect speech acts. In the limiting case what the sentence means and what the speaker means may be exactly the same; for example, the speaker might in a certain context utter the sentence "The cat is on the mat" and mean exactly and literally that the cat is on the mat. Strictly speaking, the expression 'literal' in the phrase 'literal meaning of the sentence' is pleonastic since all these other sorts of meaning – ironical meaning, metaphorical meaning, indirect speech acts and conversational implications – are not properties of sentences at all, but rather of speakers' utterances of sentences.

For sentences in the indicative, the meaning of the sentence determines a set of truth conditions; that is, it determines a set of conditions such that the literal utterance of the sentence to make a statement will be the making of a true statement if and only if those conditions are satisfied. According to some accounts, to know the meaning of such a sentence is just to know its truth conditions. Sometimes the meaning of a sentence is such that its truth conditions will vary systematically with the contexts of its literal utterance. Thus the sentence "I am hungry" might be uttered by one person on one occasion to make a true statement and yet be uttered by another person, or by the same person on another occasion, to make a false statement. Such 'indexical' or 'token reflexive' sentences differ from sentences such as "Snow is white" whose truth conditions do not vary with the context of utterance. It is important to notice however that the notion of the meaning of a sentence is absolutely context free. Even in the case of indexical sentences the meaning does not change from context to context; rather the constant meaning is such that it determines a set of truth conditions only relative to a context of utterance. The literal meaning of the sentence is the meaning it has independently of any context whatever; and, diachronic changes apart, it keeps that meaning in any context in which it is uttered.

Something like the picture sketched above provides a set of assumptions behind recent discussions in 'semantics' and 'pragmatics' that is so per-

vative as hardly to constitute a theory at all; it is rather the framework within which any theory must be stated and validated. True, there have been skeptical doubts about various aspects of it. Some philosophers have argued that this notion of meaning is insufficiently empirical and should be replaced by some more behavioristic surrogates stated in terms of the stimulus and response patterns of speakers and hearers. Some have argued that the picture leads to an unwarranted hypostatization of meanings as separate entities. I think both of these objections are invalid, but I will not argue the points here. Furthermore there are some variations on this received opinion which contain fairly serious mistakes, and I will mention one such if only to get it out of the way at the beginning. Some philosophers and linguists mistakenly suppose that the distinction between sentence and utterance is the same as the distinction between type and token, and that utterances just are identical with sentence tokens. They then suppose that because utterance meaning can differ from sentence meaning that somehow sentence tokens acquire 'different meanings in context' from the meaning of sentence types, which are context free. I believe that both aspects of this view are mistaken. That is, it is a category mistake to suppose that an utterance of a token and a token are identical and it is a mistake (derived from the previous one) to suppose that where utterance meaning differs from sentence meaning, the token acquires a different meaning from the type. If an argument is needed to show that these are mistakes it should be sufficient to point out that an utterance could not be identical with a token, because the same utterance can involve many tokens, as when one publishes one's utterances in printed form, and the same token can be used in making of several utterances as for example when one holds up the same 'STOP' sign on several occasions. Every utterance does indeed involve the production or use of a token, but the utterance is not identical with the token, and where utterance meaning differs from sentence meaning, the token does not change its meaning. Barring diachronic changes, special codes, and the like, the meaning of the token is always the same as the meaning of the type. Sentence meaning, type or token, needs to be distinguished from the speaker's utterance meaning, and the sentence-utterance distinction is not the same as the type-token distinction.

Ignoring then the skeptical doubts about meaning that have been expressed about the received opinion, and setting aside those versions of it

that contain definite mistakes, I shall argue in what follows that while the received opinion is for the most part correct (in particular it is correct in emphasizing the distinction between sentence meaning and utterance meaning), it errs in presenting the notion of the literal meaning of the sentence as a context free notion. Rather, I shall argue that for a large number of cases the notion of the literal meaning of a sentence only has application relative to a set of background assumptions, and furthermore these background assumptions are not all and could not all be realized in the semantic structure of the sentence in the way that presuppositions and indexically dependent elements of the sentence's truth conditions are realized in the semantic structure of the sentence.

My strategy in constructing the argument will be to consider sentences which appear to be favorable cases for the view that literal meaning is context free and then show that in each case the application of the notion of the literal meaning of the sentence is always relative to a set of contextual assumptions. Consider the sentence "The cat is on the mat". If any sentence has a clear literal meaning independent of any context this old philosophical chestnut ought to be it. To be sure, it contains indexical elements. In understanding an utterance of the sentence to make a statement we need to know which cat and which mat are being referred to and at what time and place the cat is being said to be on the mat. But these context dependent features of presupposition and indexicality are already realized in the semantic elements of the sentence, and if they are unclear in any particular utterance we could always make them more explicit by adding more indexical elements to the sentence – this cat right here is now on this mat right here – or we could eliminate the explicitly indexical features and substitute descriptions and time and space coordinates – the cat which has such and such features is on the mat with so and so features at such and such time and place. In addition to its indexical features the sentence carries a constant and unvarying descriptive meaning which the indexical elements serve to nail down to specific contexts in specific utterances. This unvarying descriptive content determines the truth conditions of the sentence, which the indexical elements relate to specific contexts of utterance. We might, with apologies for poor draftsmanship, represent this descriptive element as follows (see Figure 1, page 211).

When things are like that, we feel inclined to say, the cat is on the mat; otherwise not. And that is what the sentence says – it says things in



Fig. 1

the cat and mat line of business are in the relation depicted. Of course, we might concede, the sentence is not as determinate as the picture, for the cat might be sitting or standing on the mat or facing the other way and still the truth conditions of the sentence would be satisfied; and we might also concede that there is the problem of vagueness. If the cat was half on and half off the mat we might not know what to say, but such concessions raise no difficulty with our notion of context independent literal meaning.

But now suppose that the cat and the mat are in exactly the relations depicted only they are both floating freely in outer space, perhaps outside the Milky Way galaxy altogether. In such a situation the scene would be just as well depicted if we turned the paper on edge or upside down since there is no gravitational field relative to which one is above the other. Is the cat still on the mat? And was the earth's gravitational field one of the things depicted in our drawing?

What I think it is correct to say as a first approximation in answer to these questions is that the notion of the literal meaning of the sentence "The cat is on the mat" does not have a clear application, unless we make some further assumptions, in the case of cats and mats floating freely in outer space; and though our picture did not depict the earth's gravitational field, it like the sentence only has an application relative to a set of background assumptions.

Well, it might be said in response to this, if these are really assumptions behind the notion of the literal meaning of the sentence, why not make them perfectly clear as further truth conditions of the sentence? They could be treated as further Strawsonian presuppositions, or if we did not want to treat them as truth conditions, they could be stage directions for the applicability of the sentence. That is, just as, according to some philosophers, 'is bald' in the sentence "The king of France is bald" only has application if there is a king of France, so we might say the descriptive

meaning of the sentence "The cat is on the mat" only has application at or near the surface of the earth or in some other similar gravitational field. But that, like any other presupposition, can be made explicit as part of the meaning of the sentence. What the sentence really means is expressed by: "(At or near the surface of the earth or some similar gravitational field) the cat is on the mat". Or alternatively we could treat this condition as a further stage direction for the application of the sentence, but still the stage directions would be a part of the literal meaning, at least in the sense that they would be made completely explicit in the semantic analysis of the sentence. On this account the sentence is rendered: "The cat is on the mat (this sentence only applies at or near the surface of the earth or in some similar gravitational field)". But these answers to our difficulty won't do, for at least two reasons. First, it is not always the case that the literal application of the sentence requires a gravitational field. That is, it is easy to construct examples where it would be quite literally true to say that the cat is on the mat, even though there is no gravitational field. For example, as we are strapped in the seats of our space ship in outer space we see a series of cat-mat pairs floating past our window. Oddly, they come in only two attitudes. From our point of view they are either as depicted in figure 1, or as would be depicted if figure 1 were upside down. "Which is it now?", I ask. "The cat is on the mat", you answer. Have you not said exactly and literally what you meant?

But secondly, even if we got all these assumptions about gravitational fields somehow represented as part of the semantic content of the sentence, we would still be left with an indefinite number of other contextual assumptions that we would have to deal with. Consider the following example. Suppose the cat and the mat are in the spatial relations depicted in figure 1, at the surface of the earth, but that each, cat and mat, are suspended on an intricate series of invisible wires so that the cat, though slightly in contact with the mat, exerts no pressure on it. Is the cat still on the mat? Once again it seems to me that the question does not have a clear answer, and that is just another way of saying that the meaning of the sentence "The cat is on the mat" does not have a clear application in the context as so far specified and hence it does not yet determine a clear set of truth conditions. And once again, it seems to me we can easily fill in the context to give the sentence a clear application. Suppose the cat and the mat are

part of a stage set. The wires are there to facilitate rapid movement of the props, as the cat has to be moved from chair to mat to table. "Where is it now?" the director shouts from backstage; "The cat is on the mat" shouts his assistant. Does he not say exactly and literally what he means? Further examples of the contextual dependence of the applicability of the literal meaning of this sentence are easy to generate. Suppose that the mat is as stiff as a board and is stuck into the floor at an angle. Suppose the cat is drugged into a stupor and is placed relative to the mat in the following attitude.




Fig. 2

Does this situation satisfy the truth conditions of "The cat is on the mat?" Again, I feel inclined to say that so far the question does not have a clear answer, and that relative to one set of additional assumptions the situation does satisfy the truth conditions of the sentence, relative to another set it does not; but this variation has nothing to do with vagueness, indexicality, presupposition, ambiguity, or any of the other stocks in trade of contemporary 'semantic' and 'pragmatic' theory, as these notions are traditionally conceived. Suppose the cat's owner is in the next room, while I unbeknownst to him have drugged his cat and stiffened his mat with my special stiffening solution. "Where is the cat?" asks the owner from his position next door. "The cat is on the mat", I answer. Have I told the truth? My inclination is to say that my answer is misleading at best and probably should be described as an ingenious lie, since I know that that is not what the owner understands when he hears and gives a literal interpretation to the utterance of the sentence, "The cat is on the mat". But now consider a different variation of the same example. The mat is in its stiff angled position, as in figure 2, and it is part of a row of objects similarly sticking up at odd angles – a board, a fence post, an iron rod, etc. These facts are known to both speaker and hearer. The cat

jumps from one of these objects to another. It is pretty obvious what the correct answer to the question, "Where is the cat?" should be when the cat is in the attitude depicted in figure 2: The cat is on the mat.

These examples are designed to cast doubt on the following thesis: Every unambiguous sentence,² such as "The cat is on the mat" has a literal meaning which is absolutely context free and which determines for every context whether or not an utterance of that sentence in that context is literally true or false.

The examples are further designed to support the following alternative hypothesis: For a large class of unambiguous sentences such as "The cat is on the mat", the notion of the literal meaning of the sentence only has application relative to a set of background assumptions. The truth conditions of the sentence will vary with variations in these background assumptions; and given the absence or presence of some background assumptions the sentence does not have determinate truth conditions. These variations have nothing to do with indexicality, change of meaning, ambiguity, conversational implicature, vagueness or presupposition as these notions are standardly discussed in the philosophical and linguistic literature.

Perhaps the thesis that literal meaning is absolutely context free could be replaced by a weaker thesis: while there may indeed be a large class of sentences whose literal meaning only determines a set of truth conditions relative to a set of background assumptions, still (it might be argued) for each sentence in this class we can specify these assumptions in such a way that they will be constant for every literal occurrence of the sentence.

But our examples have already cast doubt even on this weaker thesis; for the truth conditions of the sentence "The cat is on the mat" are satisfied in each of our 'abnormal' contexts provided that the abnormal context is supplemented with some other assumptions. And thus, there is no constant set of assumptions that determine the applicability of the notion of literal meaning, rather the sentence may determine different truth conditions relative to different assumptions in ways that have nothing to do with ambiguity, indexical dependence on context, presupposition failure, vagueness or change of meaning as these notions are traditionally conceived. Furthermore, our examples suggest that the assumptions are not specifiable as part of the semantic content of the sentence, or as presuppositions of the applicability of that semantic content, for at least two reasons. First, they are not fixed and definite

in number and content; we would never know when to stop in our specifications. And second, each specification of an assumption tends to bring in other assumptions, those that determine the applicability of the literal meaning of the sentence used in the specification. It is important to note, however, that so far my examples only challenge the idea that there is a sentence by sentence specification of the background assumptions as part of the semantic analysis of each sentence; I have not so far addressed the question whether it might be possible to give a specification of all the assumptions against which speakers understand and apply the literal meanings of sentences.

Notice that we used only very limited resources in constructing the examples. We concentrated only on the contextual dependence of the word 'on' as it occurs in the sentence. If we went to work on 'cat' or 'mat' we could find much more radical forms of contextual dependence. Second, we imagined no changes in the laws of nature. Again, given freedom to mess around with the laws of nature, I believe we could get still more radical breakdowns in the view that the application of the literal meaning of sentences is absolutely context independent.

II

Corresponding to the notion of the truth conditions of an indicative sentence is the notion of the obedience conditions of an imperative sentence or the fulfillment conditions of an optative sentence, and many of these same points will emerge if we consider sentences in the imperative and other moods. Suppose I go into a restaurant determined to say exactly and literally what I mean, that is determined to utter imperative sentences that give exact expression to my desires. I start by saying: "Give me a hamburger, medium rare, with ketchup and mustard, but easy on the relish". I will remark first of all that a prodigious amount of background information has already been invoked even by the example as so far described – entire institutions of restaurants and money and of exchanging prepared foods for money, for a start; and it is hard to see how the sentence could have quite the same obedience conditions if these institutions did not exist, or if the same sentence were uttered in a radically different context, if for example the sentence were uttered by a priest as part of a prayer or tacked onto the end of his inaugural swearing in by

an incoming President of the U.S. Still, one might argue, the sense of 'give' in which it initiates commercial transactions is in part defined by these systems of constitutive rules. So this much of contextual dependence is in part realized in the semantic structure of the sentence. But even if one conceded that – and it is not clear that one should – there are all sorts of other assumptions on which the application of the sentence rests and which are not even remotely close to being realized in the semantic structure of the sentence. Suppose for example that the hamburger is brought to me encased in a cubic yard of solid lucite plastic so rigid that it takes a jack hammer to bust it open,³ or suppose the hamburger is a mile wide and is 'delivered' to me by smashing down the wall of the restaurant and sliding the edge of it in. Has my order "Give me a hamburger, medium rare, with ketchup and mustard, but easy on the relish" been fulfilled or obeyed in these cases? My inclination is to say no, it has not been fulfilled or obeyed because that is not what I meant in my literal utterance of the sentence (though again it is easy to imagine variations in our background assumptions where we would say that the order has been obeyed). But the fact that the order has not been obeyed – that is, that the obedience conditions of the sentence are not satisfied relative to that context – does not show that I failed to say exactly and literally what I meant, that what I should have said is "Give me a hamburger, medium rare, with ketchup and mustard, but easy on the relish; and don't encase it in plastic and no mile wide hamburgers, please". If we say that, then it will become impossible ever to say what we mean because there will always be further possible breakdowns in our background assumptions which would lead us to say that the obedience conditions of the sentence were not satisfied in a given context. Rather it seems to me what we should say in such cases is that I did say exactly and literally what I meant but that the literal meaning of my sentence, and hence of my literal utterance, only has application relative to a set of background assumptions which are not and for the most part could not be realized in the semantic structure of the sentence. And there are, in these cases as in the indicative cases, two reasons why these extra assumptions could not all be realized in the semantic structure of the sentence, first they are indefinite in number, and second, whenever one is given a literal statement of these assumptions, the statement relies on other assumptions for its intelligibility.

Examples of the contextual dependence of the application of the notion of literal meaning are easy to multiply. Consider the imperative sentence "Shut the door". As soon as we hear this sentence we are likely to picture a standard scene in which it would have a clear literal application. The speaker and hearer are in a room. The room has an open door that can be moved into its door frame and latches into the frame when closed. But as soon as we alter this domestic scene radically the sentence loses its application. Suppose the speaker and hearer are floating with a door in the middle of the ocean; or suppose they and the door are sitting alone in the Sahara. What are the obedience conditions of "Shut the door" in these situations? Still, one might say, though the literal meaning of the sentence loses its application in these mid-ocean and Sahara examples, such presuppositions as that there is a room and that the door is in the door frame do look like standard presuppositions for the obedience conditions of this sentence. They at least can be spelled out. But in answer to that: when we have spelled them out we are no better off. Suppose we are in our standard scene and the speaker utters the sentence "Shut the door", saying exactly and literally what he means. Suppose the hearer goes to the door and chops the entire complex – door, frame, hinges, latch and all – from the wall, sets the whole mess up in the middle of the room and then moves the door on its hinges so that it latches in the frame. Has he shut the door, that is, are the obedience conditions of the sentence satisfied? I am inclined to say that, as we look over our shoulder at the gaping hole he left in the wall, we would say no, the obedience conditions are not satisfied. But again, it would be very easy to vary the assumptions in such a way that we would say the obedience conditions of the sentence were satisfied. And there does not seem to be any upper limit on our ability to generate such deviant contexts. Suppose the hearer swallows the whole thing – wall, door, frame, and latch – and then moves the door into the door frame as part of the peristaltic contraction of his gut during digestion. Did he "shut the door"? Furthermore, the sentences that we used to state the presuppositions – that there is a room, that the room has walls, that at least one wall has a door opening with a door frame – will be just as subject to the sorts of contextual dependencies we designed them to eliminate as was our original sentence. Yet, for each of these examples, I want to say that the speaker says or can say exactly and literally what he means. There is no question of his being ambiguous, vague, or meta-

phorical when he says "Give me a hamburger..." or "Shut the door"; but these literal utterances only determine a set of obedience conditions relative to a set of contextual assumptions. Different assumptions may determine different obedience conditions; and for some assumptions there may be no obedience conditions at all, even though, to repeat, the sentence and hence the utterance is perfectly unambiguous.

Indeed, the very terminology of 'assumptions' and 'contexts' might mislead if it suggests that for each sentence we could make all these assumptions explicit on the model of a set of axioms, as Peano and Frege tried to make explicit the assumptions of arithmetic, or as an economic theorist constructing a deductive economic model makes explicit his assumptions in the form of a set of axioms. But even assuming we could not do a sentence by sentence specification of the assumptions behind the understanding and application of each sentence, could we do a completely general specification of all the assumptions, all the things we take for granted, in our understanding of language? Could we make our whole mode of sensibility fully explicit? It seems to me that the arguments in this article don't determine the answer to that question one way or the other. The fact that for each of a large range of sentences the assumptions are variable and indefinite and that the specification of one will tend to bring in others does not by itself show that we could not specify an entire set which would be independent of the semantic analysis of individual sentences but which taken together would enable us to apply the literal meaning of sentences. The practical difficulties in any such specification would of course be prodigious, but is there any theoretical obstacle to the task? In order to show that there was we would have to show that the conditions under which sentences can represent were not themselves fully representable by sentences. Perhaps that claim is true but it has not been the aim of my discussion to show that it is true.

The claims that I have made about sentences lead naturally to our next conclusion: what I have said about literal meaning also applies to intentional states⁴ in general. A man who believes that the cat is on the mat or who expects that they will bring him a medium rare hamburger has his belief and expectation only against a background of other inexplicit 'assumptions'. Just as the literal meaning of a sentence will determine different truth or obedience conditions relative to different sets of assumptions, so a belief or expectation will have different conditions of satis-

faction relative to different sets of assumptions. And it is really not surprising that there should be this parallelism between literal meaning and intentional states, since the notion of the literal meaning of a sentence is in a sense the notion of conventional and hence fungible intentionality: it is what enables the sentence to represent out there in public, so to speak; whereas my beliefs, desires, and expectations just represent their conditions of satisfaction *tout court*, regardless of whether they get any help from having public forms of expression. The general point is that representation, whether linguistic or otherwise in general goes on against a background of assumptions which are not and in most cases could not also be completely represented as part of or as presuppositions of the representation, for the two reasons we have already stated: the assumptions are indefinite in number and any attempt to represent them will tend to bring in other assumptions. There is an obvious analogy with pictorial representation in this last consideration, because if one tries to depict the method of projection of one's picture in yet another picture, the second picture will also require an as yet undepicted method of projection.

III

It is important not to overstate the case that has been made so far. I have by no means demonstrated the contextual dependence of the applicability of the notion of the literal meaning of a sentence. Rather, I have offered a few examples together with some hints as to how we could generalize the phenomena discovered in those examples. Furthermore since the examples concern weird cases it is hard to be sure about our linguistic intuitions in describing them. But even assuming I am right about these examples perhaps we could find sentences for which there would be no such contextual dependency. Perhaps one might show, for example, that an arithmetical sentence such as " $3 + 4 = 7$ " is not dependent on any contextual assumptions for the applicability of its literal meaning. Even here, however, it appears that certain assumptions about the nature of mathematical operations such as addition⁵ must be made in order to apply the literal meaning of the sentence. But to this one might reply, in a logicist vein, that these assumptions are in a sense part of meaning of the sentence. Such an argument would raise many of the traditional disputes

in the philosophy of mathematics, and I shall not attempt to pursue them here. For the purposes of this discussion, it is sufficient to argue that the notion of absolutely context free literal meaning does not have general application to sentences; and indeed there does seem to be a very large class of sentences to which we could extend the sorts of arguments for the contextual dependence of the applicability of literal meaning that we discussed earlier.

There are two skeptical conclusions that these reflections might seem to imply that I want to renounce explicitly. First, I am not saying that sentences do not have literal meanings. To show that a phenomenon *X* can only be identified relative to another phenomenon *Y* does not show that *X* does not exist. To take an obvious analogy, when one says that the notion of the movement of a body only has application relative to some coordinate system, one is not denying the existence of motion. Motion, though relative, is still motion. Similarly, when I say that the literal meaning of a sentence only has application relative to the coordinate system of our background assumptions, I am not denying that sentences have literal meanings. Literal meaning, though relative, is still literal meaning.

Well then, what is meant by 'application' when I say that the literal meaning of a sentence only has 'application' relative to a set of background assumptions? Simply this. There are certain jobs that we want the notion of meaning to do for us; it connects in all sorts of systematic ways with our theory of language and with our pretheoretical beliefs about language. Meaning is tied to our notions of truth conditions, entailment, inconsistency, understanding, and a host of other semantic and mental notions. Now the thesis of the relativity of meaning is the thesis that one can only make these connections relative to some coordinate system of background assumptions. In the case, for example, of truth conditions (or obedience conditions for imperative sentences) the thesis of the relativity of meaning has the consequence that the sentence may determine one set of truth conditions relative to one set of assumptions and another set relative to another set of assumptions even though the sentence is not ambiguous and the variation is not a matter of indexical dependence (analogously, the thesis of the relativity of motion has the consequence that the same object at the same time may be moving one way relative to one coordinate system and the other way relative to another coordinate system even though it is not moving in two different ways), and without some set of background

assumptions the sentence does not determine a definite set of truth conditions at all. For most sentences of the “Cat is on the mat”, “Bill is in the kitchen”, “My car has a flat tire” variety, the background assumptions are so fundamental and so pervasive that we don’t see them at all. It takes a conscious effort to prise them off and examine them, and incidentally, when one does prise them off it tends to produce an enormous sense of annoyance and insecurity in philosophers, linguists, and psychologists – or at any rate such has been my experience.

A second skeptical conclusion that I explicitly renounce is that the thesis of the relativity of literal meaning destroys or is in some way inconsistent with the system of distinctions presented in my brief summary of the standard account of meaning at the beginning of this article, namely that system of distinctions that centers around the distinction between the literal sentence meaning and the speaker’s utterance meaning, where the utterance meaning may depart in various ways from literal sentence meaning. The distinction, for example, between literal sentence meaning and metaphorical or ironical utterance meaning remains intact. Similarly, the distinction between direct and indirect speech acts remains intact. The modification that the thesis of the relativity of meaning forces on that system of distinctions is that in the account of how context plays a role in the production and comprehension of metaphorical utterances, indirect speech acts, ironical utterances, and conversational implications, we will need to distinguish the special role of the context of utterance in these cases from the role that background assumptions play in the interpretation of literal meaning. Furthermore there is nothing in the thesis of the relativity of literal meaning which is inconsistent with the Principle of Expressibility, the principle that whatever can be meant can be said. It is not part of, not a consequence of, my argument for the relativity of literal meaning that there are meanings that are inherently inexpressible.

In the face of these examples to support the thesis of the relativity of literal meaning, defenders of the traditional theory of absolute literal meaning are likely to resort to certain standard moves and it is perhaps well to obviate these before they can even get started. Neither the sentence-utterance distinction (much less the type-token distinction) nor the performance-competence distinction will rescue the thesis of absolutely context free literal meaning as far as our counterexamples are concerned. The discussion throughout has been about sentences, and I have discussed

the meaning of utterances only in cases where utterance meaning coincides with sentence meaning, that is, only in cases where the speaker means literally what he says. Furthermore I have been discussing the understanding of the literal meaning of a sentence by a speaker as part of the speaker's semantic competence. The thesis I have been advancing is that for a large class of sentences the speaker, as part of his linguistic competence, knows how to apply the literal meaning of a sentence only against a background of other assumptions. If I am right, this argument has the consequence that there is no sharp distinction between a speaker's linguistic competence and his knowledge of the world, but there are numerous other arguments to support that position anyway. I have frequently made use of arguments of the form "What would we say if . . ." but that does not mean that we are not discussing sentence meanings or that we are not discussing linguistic competence.

Why should things be the way I have described them, assuming for the sake of argument that my description is correct? That is, for example, why couldn't we just lay it down that the meaning of the sentence "The cat is on the mat" or "The door is shut" was going to be absolutely context free? Meanings are, after all, a matter of convention, and if heretofore such conventions have rested on background assumptions why not put an end to this dependence by a new convention that there shall henceforth be no such dependence? I don't know how to answer these questions except by saying that literal meaning is dependent on context in the same way that other non-conventional forms of intentionality are dependent on context, and there is no way to eliminate the dependence in the case of literal meaning which would not break the connections with other forms of intentionality and hence would eliminate the intentionality of literal meaning altogether.

Since perception is in all likelihood the primary form of intentionality, the one on which all others depend, we can best begin by showing the contextual dependency of the applicability of the contents of our perceptions. Consider the characteristic visual experiences that would be present when we are in a position to say "I see that the cat is on the mat". As far as the purely qualitative visual aspects of these experiences are concerned (and I don't know a better vocabulary than 'qualitative visual aspects' to get at what I am talking about) many of these aspects could have been produced by any number of causes and in any number of

situations. They might have been produced by stimulating the optical centers of my brain in such a way as to give me experiences with visual aspects just like the aspects I have in my present visual experiences. Yet I want to say that my present visual experiences, the ones that enable me to say that I see that the cat is on the mat, have a form of intentionality that these other experiences would not have, assuming that is that I knew what was going on in the two cases. In my present experience I assume that I am perceiving the cat and the mat from a certain point of view where my body is located; I assume that these visual experiences are causally dependent on the state of affairs that I perceive; I assume that I am not standing on my head and seeing cat and mat upside down, etc.; and all these assumptions are in addition to such general assumptions as that I am in a gravitational field, there are no wires attaching to cat and mat, etc. Now, the intentionality of the visual experience will determine a set of conditions of satisfaction. But the purely visual aspects of the experience will produce a set of conditions of satisfaction only against a set of background assumptions which are not themselves part of the visual experience. I don't for example see the point of view from which I see that the cat is on the mat and I don't see the gravitational field within which they are both located. Yet the conditions of satisfaction which are determined by the content of my perception are in part dependent on such assumptions. Indeed in this case as in the literal meaning case, the intentionality of the visual perception only has an application, only determines a set of conditions of satisfaction, against some system of background assumptions. Thus there seems no way to eliminate the contextual dependence of literal meaning since it is built into other forms of intentionality on which literal meaning depends. To borrow an expression from Wittgenstein, it is part of the grammar of "The cat is on the mat" that this is what we call "seeing that the cat is on the mat", "believing that the cat is on the mat", etc. There is no way to eliminate the contextual dependence of the sentence "The cat is on the mat" without breaking the connections between that sentence and the *perception* that the cat is on the mat, or the belief that the cat is on the mat, and it is on such connections that the meaning of the sentence depends.

NOTES


¹ I am indebted to many people for discussion of the ideas in this paper, and I especially want to thank M. Bierwisch, H. Dreyfus, and D. Searle.

² One can of course locate ambiguities even in this sentence; e.g., 'cat' is sometimes used as a slang expression for caterpillar tractors. But such ambiguities are irrelevant to our present discussion.

³ This example was originally suggested to me by H. Dreyfus in a discussion of another issue.

⁴ By 'intentional states' here I mean those mental states such as belief and desire that are directed at or about objects and states of affairs in the world. They differ from such states as pains and tickles that are not in that way directed at or about anything.

⁵ Thus in Wittgenstein's example $A = 3$, $B = 4$, but $A + B = 5$



Manuscript received 20 November 1977