

MEANING

BY

Stephen R. Schiffer



SCHIFFER

Chapter 2 of

Meaning

OXFORD

AT THE CLARENDON PRESS

1972

once agents have these propositional attitudes, they will communicate with one another. Once agents begin to communicate with one another they will begin to develop a conventional system of communication. Once even a rudimentary "language" or conventional system of communication is possessed by a group of agents it will then become possible for them to have propositional attitudes which they could not otherwise have; and this will make it possible for them to communicate things which they could not otherwise communicate, which in turn will result in a more sophisticated "language", which in turn will make it possible for them to have propositional attitudes they could not otherwise have, and so on.

On this view, then, the fact (if it is a fact) that there are certain propositional attitudes one cannot have unless one has a language is no objection to a Gricean approach to meaning and language. Certain aspects of this view will be elaborated in chapter V.

II

TOWARD AN ACCOUNT OF
S-MEANING*Introduction*

ONE way of finding out what meaning is is to see what meaning is not, and in this chapter I will be mainly concerned to discuss certain objections to Grice's account of *S*-meaning; objections, first, to the alleged sufficiency of Grice's conditions, and objections, lastly, to the alleged necessity of one of Grice's conditions. In between I will try to show, in part, what must be added to Grice's account in order to arrive at a set of conditions which are jointly sufficient for someone's meaning something by uttering *x*. The objections I discuss in this chapter do not exhaust those that can and have been made against Grice's account of *S*-meaning. Some of these objections will be dealt with in subsequent chapters.

I shall begin with the restatement of Grice's account given on pp. 12-13.

II.1 *Some objections to the alleged sufficiency of Grice's analysis*

(a) Here is a slightly more detailed version of a counter-example of Strawson's.¹

S wants to get *A* to believe that the house *A* is thinking of buying is rat-infested. *S* decides to bring about this belief in *A* by taking into the house and letting loose a big fat sewer rat. For *S* has the following scheme. He knows that *A* is watching him and knows that *A* believes that *S* is unaware that he, *A*, is watching him. It is *S*'s intention that *A* should (wrongly) infer from the fact that *S* let the rat loose that *S* did so with the intention that *A* should arrive at the house, see the rat, and, taking the rat as "natural evidence", infer therefrom that the house is rat-infested. *S* further intends *A* to realize that given the nature of the rat's arrival, the existence of the rat cannot be taken as genuine or natural evidence that the house is rat-infested; but *S* knows

¹ P. F. Strawson, "Intention and Convention in Speech Acts", pp. 446-7.

that *A* will believe that *S* would not so contrive to get *A* to believe the house is rat-infested unless *S* had very good reasons for thinking that it was, and so *S* expects and intends *A* to infer that the house is rat-infested from the fact that *S* is letting the rat loose with the intention of getting *A* to believe that the house is rat-infested. Thus *S* satisfies the conditions purported to be necessary and sufficient for his meaning something by letting the rat loose: *S* lets the rat loose intending (4) *A* to think that the house is rat-infested, intending (1)–(3) *A* to infer from the fact that *S* let the rat loose that *S* did so intending *A* to think that the house is rat-infested, and intending (5) *A*'s recognition of *S*'s intention (4) to function as his reason for thinking that the house is rat-infested. But even though *S*'s action meets these conditions, Strawson is clearly right in claiming that "this is clearly not a case of attempted *communication* in the sense which . . . Grice is seeking to elucidate" (op. cit., p. 447).

What feature of this example makes it a counter-example? Strawson suggests that a minimum further condition of *S*'s trying to communicate with *A* is that "he should not only intend *A* to recognize his intention to get *A* to think that *p*, but that he should also *intend A to recognize his intention to get A to recognize his intention to get A to think that p*" (ibid.). If Strawson is right, it would appear that we must add to the revised set of conditions the further condition that *S* have the intention

(6) that *A* should recognize *S*'s intention (3).

(b) Strawson is wary of claiming that even with this addition Grice's analysis provides sufficient conditions. Strawson's caution is, I believe, vindicated by the following counter-example.

S, who has a hideous singing voice, intends (4) to bring about *A*'s leaving the room by singing "Moon Over Miami". Further, *S* intends (1)–(3) that *A* should recognize that *S* is singing "Moon Over Miami" and infer therefrom that *S* is doing this with the intention (4) of getting *A* to leave the room, and *S* also intends (6) that *A* recognize *S*'s intention (3) that *A* recognize *S*'s intention to get *A* to leave the room (for *S* wishes to show his disdain for *A*'s being in the room). Now *S* intends that *A* will believe that *S* plans to get rid of *A* by means of *S*'s repulsive singing, but *S* expects and intends (5) that *A*'s reason for leaving the room will in fact be *A*'s recognition of *S*'s intention (4) to get him to leave the room. In other words, while *A* is intended to *think* that

S intends to get rid of *A* by means of the *repulsive singing*, *A* is really intended to have as his reason for leaving the fact that *S* wants him to leave.

I think it is clear that *S* should not be said to have meant by singing "Moon Over Miami" that *A* was to leave the room; nor should we say that there was something else that *S* meant. Following Strawson, it would seem that the way to exclude this counter-example is to add the condition that *S* have the intention

(7) that *A* should recognize *S*'s intention (5).

Putting all of this together, Grice's analysans will have grown on the suggested way of revision to the following not uncomplicated set of conditions:

S meant something by (or in) uttering *x* iff *S* uttered *x* intending

- (1) that *x* have a certain feature(s) *f*;
- (2) that a certain audience *A* recognize that *x* is *f*;
- (3) that *A* infer at least in part from the fact *x* is *f* that *S* uttered *x* intending (4):
- (4) that *S*'s utterance of *x* produce a certain response *r* in *A*;
- (5) that *A*'s recognition of *S*'s intention (4) shall function as at least part of *A*'s reason for his response *r*;
- (6) that *A* recognize *S*'s intention (3);
- (7) that *A* recognize *S*'s intention (5).

In effect this restatement says that *S* meant something by uttering *x* just in case *S* uttered *x* intending *A* to recognize that *S* uttered *x* intending to produce a response *r* in *A* by means of *A*'s recognition of *S*'s intention to produce *r* in *A*; that is, not only must *S* utter *x* with the original complex Gricean intention, he must also intend *A* to recognize this.

It will be relevant to notice that there is more of a rationale for the addition of this further intention than its use to rule out a few *recherché* counter-examples. For, in general, *S* can utter *x* intending to produce a certain response *r* in *A* by means of *A*'s recognition of this intention *only if* *S* expects *A* to recognize that *S* intends to produce *r* in *A* by means of recognition of intention *or else* *S* intends to *deceive A* as to the means by which *S* intends to produce *r* in *A*. In other words, given that *S* intends to produce *r* in *A* by means of recognition of intention (in the relevant sense,

of course) and given that *S* does not want to deceive *A*, then *S* must—on pain of not satisfying his primary intention to produce *r* in *A*—expect *A* to think that *S* intends *A*'s reason (or part of his reason) for his response *r* to be the fact that *S* intends to produce *r* in *A*. I believe this can be shown to be so by the following argument. (For simplicity I use an example in which the intended response is *A*'s thinking that *p*; the argument should, *mutatis mutandis*, apply to any other type of response for which *A* has reasons.)

(i) A necessary condition of one's doing an act with the intention of thereby bringing about some further effect or result is that one should think that one's doing that act will (or might—this qualification should be made throughout) be sufficient (in the circumstances) for the production of that effect. So if *S* intends *A* to think that *S* uttered *x* intending to get *A* to think that *p*, *S* must think that *A* thinks (or will think) that *S* thinks that his utterance of *x* will be sufficient in the circumstances for getting *A* to think that *p*. (ii) Since in general the only way of getting someone to think that *p* is to provide that person with a reason (not necessarily a good reason but something which will be taken as a good reason) for thinking that *p*, *A* will not, in general, think that *S* uttered *x* intending to get *A* to think that *p* unless *A* thinks that *S* thinks that his utterance *x* (or some product thereof—this qualification can be ignored) will be taken by *A* to be a reason for thinking that *p*. Consequently, given that *S* intends *A* to think that *S* uttered *x* intending to get *A* to think that *p*, *S* must think that *A* thinks (or will think) that *S* thinks that *A* will take *S*'s utterance *x* as a reason for thinking that *p*. (For example, Herod presents Salome with the severed head of St. John the Baptist intending to inform her that St. John the Baptist is dead and also intending her to think that he intended to inform her that St. John the Baptist is dead. It was possible for Herod to have the intention to inform Salome that St. John the Baptist is dead because he knew that she knew that one could not live without one's head and because he knew that she would recognize the head on the charger as the head of St. John the Baptist. It was possible for Herod to intend Salome to recognize his first intention because he knew that she knew that he knew that she knew that one could not live without one's head and because he knew that she knew that he knew that she would recognize the head on the charger as being the head of St. John

the Baptist.) (iii) In general, *S* will think that *A* will think that the fact that *S* uttered *x* is a reason for thinking that *p* only if *S* thinks that *A* has a certain belief(s) *r* which would warrant *A*'s thinking that *p* on the basis (in part) of thinking that *S* uttered *x*. Strictly, it is only necessary that *S* think *A* has some belief or beliefs which would warrant his thinking that *p* on the basis of thinking that *S* uttered *x*. But ordinarily *S* will be justified in thinking *A* has some such belief only if *S* thinks that *A* has a certain belief(s) which would warrant *A*'s thinking that *p* on the basis of thinking *S* uttered *x*. Since this is so, *A* will, in general, think that *S* uttered *x* intending thereby to provide *A* with a reason for thinking that *p* only if *A* thinks that there is a certain belief(s) such that *S* thinks both that *A* has (or will have) that belief(s) and that *A* thinks (or will think) that that belief(s) would warrant *A*'s thinking that *p* on the basis of thinking that *S* uttered *x*. So if *S* utters *x* intending to get *A* to think that *p*, he will expect *A* to have a certain belief(s) *r* which will warrant *A*'s thinking that *p* on the basis of thinking that *S* uttered *x*; and if *S* also intends *A* to think that *S* uttered *x* intending to get *A* to think that *p*, then *S* will expect *A* to think that *S* thinks that *A* has (or will have) a certain belief(s) *r*' (which may or may not be identical with *r*) which will warrant *A*'s thinking that *p* on the basis of thinking *S* uttered *x*. There are, then, *in principle*, two possibilities open to *S* if he utters *x* intending to produce in *A* the belief that *p* by means of *A*'s recognition of this intention: (1) *S* will expect *A* to think that *S* intends part of *A*'s reason for thinking that *p* to be the fact that *S* uttered *x* intending to get *A* to think that *p*, or (2) *S* will expect *A* to think that *S* thinks that *A* will have some reason other than recognition of *S*'s intention to get *A* to think that *p* as his reason for thinking that *p* on the basis of thinking *S* uttered *x*. In this second case there is an intended discrepancy between the warrant *A* is intended to have for thinking that *p* on the basis of thinking *S* uttered *x* and the warrant he is intended to think *S* expects him to have; i.e., in the second case *S* intends *A* to be deceived as to the means by which *S* intends to produce in *A* the belief that *p*. This is precisely the situation we have in each of the above two counter-examples (except, of course, that in the "Moon Over Miami" example the intended response is a practical one).

(c) Unfortunately, however, there is still room for a more subtle type of deception, and, as if conditions (1)-(7) were not

torturous enough, here is a further counter-example, a variation on the last one, to show that we still do not have a set of jointly sufficient conditions.

This time *S* sings "Tipperary" with the intention (4) of getting *A* to leave the room; the intention (1)–(3) that *A* recognize *S*'s intention (4) as a result of *A*'s inference from the fact that *S* is singing "Tipperary"; and the intention (6) that *A* recognize *S*'s intention (3) (for he again wishes to show his disdain for *A*'s being in the room). *S* also intends (5) that *A*'s reason for leaving the room will be his recognition of *S*'s intention (4). However, at this point *S* has the following plan. He intends that *A* should (wrongly) think that *S* intends *A* to think that *S* intends to get rid of *A* by means of *S*'s repulsive singing, but that *S* really intends to get rid of *A* by means of *A*'s recognition of *S*'s intention (4) to get *A* to leave. In other words, *S* intends *A* to reason as follows: "*S* intends me to think that he intends to get rid of me by means of his repulsive singing, but I recognize that he really intends my reason for leaving to be the fact that he wants to get rid of me and not, as he would falsely have me think, the fact that I can't stand his singing."

In the "Moon Over Miami" example there was an intended discrepancy between the reason *A* was intended to have for leaving and the reason *A* was intended to think he was intended to have. In the "Tipperary" example, *A* is intended to *think* (wrongly) that there is such an intended discrepancy. One might say that here there is an intended discrepancy between the reason *A* was intended to think he was intended to have for leaving and the reason he was intended to think he was intended to think he was intended to have! (One might, as Grice has, question the possibility of this case being realized.² For how could *S* reasonably expect such a complicated inference to be made from the fact that he is singing "Tipperary"? No doubt such cases require somewhat special circumstances, but this is no objection to their being counter-examples. The "special circumstances" appropriate for the "Tipperary" example might be this: after singing "Moon Over Miami" *S* has a third person, *B*, tell *A* *S*'s scheme in singing "Moon Over Miami", *B* pretending all the while to be secretly telling this behind *S*'s back. Thus, when *S* sings "Tipperary" *A* infers, *as intended*, that *S* is singing "Tipperary" with the same intentions he had in singing "Moon Over Miami".)

² H. P. Grice, "Utterer's Meaning and Intentions", p. 158.

It is clear that *S*'s scheme entails that *S* intend (7) that *A* recognize *S*'s intention (5), but I think that it is almost as clear that we should not want to say that *S* meant something by singing "Tipperary". Prima facie it might seem that the way to eliminate this type of counter-example would be to add a condition requiring *S* to intend

(8) that *A* recognize *S*'s intention (7);

but *in principle* we could keep on constructing counter-examples of the above kind, each time requiring us to add a condition of the above nature. If we grant that the only way of eliminating such counter-examples is to keep adding conditions which require *S* to have some (*n*)th-order intention that *A* recognize *S*'s (*n*–1)th-order intention, then this would seem to justify the fear expressed by Strawson that the analysis of *S*-meaning may involve an infinitely or indefinitely regressive series of intentions that intentions should be recognized. One who thinks that there are Gricean acts of communication and that the analysis of such acts does involve a regressive series of intentions will also think that this regress must be harmless. Such considerations are not likely to move a non-believer.

Grice has argued that there is no infinitely or indefinitely regressive series of intentions that intentions be recognized, harmless or otherwise.³ For, Grice suggests, there will come a point in the purported regress, i.e. some purported (*n*)th-order intention that *A* recognize *S*'s (*n*–1)th-order intention, such that it will be known by *S* to be impracticably difficult for *A* to infer from the fact that *S* uttered *x* that *S* uttered *x* with an (*n*–1)th-order intention; and since one cannot, in general, intend to bring about a result one knows one cannot bring about, it follows that there will be some (*n*–1)th-order intention in the alleged regress such that *S* cannot have a further (*n*)th-order intention that *A* recognize *S*'s (*n*–1)th-order intention. The reason Grice offers for thinking that the alleged regress must reach a point beyond which ordinary, non-supersubtle mortals cannot proceed is based upon the increasing complexity of the counter-examples which seem to force the addition of conditions (6)–(8). For even if one can construct a further workable counter-example of the above type which shows that *S* must intend (9) that *A* recognize *S*'s intention (8), it is clear that one will soon come to a counter-example

³ *Ibid.*

of this type where the calculations required to bring it off are simply too complicated to be made. It may not be precisely clear where this "cut-off" point lies (i.e. how many conditions requiring *S* to intend the recognition of some "lower-order" intention will be deemed necessary before the analysans is immune from a further counter-example of this type) but, Grice suggests, there will be *some* such cut-off point, and at that point one will have arrived at a set of conditions which are jointly sufficient and separately necessary for *S* meaning something by uttering x .

I have several objections to Grice's argument.

(1) Grice's reason for thinking that the alleged regress must reach a point beyond which it will be impracticably difficult to proceed is that there will come such a point in the construction of counter-examples of the above type. In other words, that there will come such a point in the construction of *counter-examples* is supposed to show that there will come such a point in *all* cases of communication.

I do not think this argument will do. Granted that the inferences and calculations required of *A* to some intention of *S* in uttering x in such complicated cases of *deception* as the above are extremely difficult. Nevertheless, it does not follow from this that the "inference" to a corresponding intention of *S* will also be difficult to make in the standard case of communication, where there is no deception and where everything is out in the open. To see that this assumption is false, one need only compare the difficulty *A* will have in inferring from the fact that *S* is singing "Tipperary" that *S* intends *A*'s reason for leaving to be that *S* intends *A* to leave with the analogous inference in a straightforward case where *S* simply tells *A* to leave by uttering the sentence 'Leave the room!'

But suppose that Grice is right and that the purported regress does not arise, *for the reasons given by Grice*, and that, therefore, there will be some (*n*)th-order intention that *A* recognize *S*'s (*n*-1)th-order intention such that *S*'s (*n*)th intention is required as a necessary condition for *S* meaning something by uttering x in order to rule out some "Tipperary"-type counter-example and such that an (*n*-1)th-order intention is the "uppermost" intention anyone could possibly hope to be recognized. We can then make two further objections.

(2) Whatever *S*'s (*n*)th intention is, it will be a *de facto* stopping point, one determined by the intelligence, ingenuity, and subtlety of actual people. But surely it is always possible to imagine two

people, a bit more intelligent and a bit more subtle than anyone else, who can succeed in bringing off a "Tipperary"-type inference requiring an (*n*+1)th-order intention to be added to the analysans. Therefore, even if Grice's argument to show that there is no regress is correct, the purported stopping point will not yield a set of logically sufficient conditions, which is what his analysans is meant to provide.

(3) Grice's argument rests on the assumption that the "cut-off" point (and, *a fortiori*, the final necessary condition of the analysans) is determined in the following way. First, determine the most complicated counter-example it is possible for any two people to bring off. In this example *S* will have some (*n*)th-order intention that *A* recognize *S*'s (*n*-1)th-order intention but no (*n*+1)th-order intention that *A* recognize *S*'s (*n*)th-order intention. To rule out this "uppermost" counter-example add a condition requiring *S* to have an (*n*+1)th-order intention that *A* recognize *S*'s (*n*)th-order intention. But, if Grice's reason for thinking the regress must reach a "cut-off" point is correct, then the (*n*-1)th-order intention of *S* in this "uppermost" counter-example will be the "uppermost" intention anyone could possibly hope to be recognized. Therefore, if we make it a necessary condition—in order to rule out this counter-example—that *S* have an (*n*+1)th-order intention that *A* recognize *S*'s (*n*)th-order intention that *A* recognize *S*'s (*n*-1)th-order intention, it will follow that a necessary condition of *S* meaning something by uttering x is that he intend *A* to recognize an intention he could not possibly expect *A* to recognize; so if we take this way out it will be impossible for anyone to mean anything. On the other hand, if we only require that *S* intend *A* to recognize *S*'s (*n*-1)th-order intention, then, not only will this necessarily fail to rule out at least one counter-example (and so not be sufficient), it will have the consequence that only the two most subtle and intelligent beings alive could mean anything, and they could only communicate with one another. (Grice has suggested that one way of avoiding this last objection would be to have the analysans "vary from case to case, depending on such things as the nature of the intended response, the circumstances in which the attempt to elicit the response is made, and the intelligence of the utterer and of the audience" (ibid., p. 159). But even Grice admits that it is doubtful whether this would be acceptable; and even if this way out were accepted, I fear it might still have the consequence that *S* will

always be required to have *some* intention which he does not think *A* could recognize and which is necessary for his meaning something in the circumstances, and so *A* could never (except, perhaps, where *S* miscalculated) know that *S* meant something.)

What makes each of the examples (a)–(c) a counter-example is that *S* intends to deceive *A* in one way or another. It might be thought that these counter-examples could be eliminated without involving a proliferation of intentions by requiring that *S* *not* have certain intentions. Grice has suggested the possibility that instead of eliminating the above counter-examples by the addition of conditions (6)–(8), we add instead to the analysis a condition requiring it to be the case that:

there is no inference-element *E* such that [*S*] uttered *x* intending both (1') that *A*'s determination of *r* should rely on *E* and (2') that *A* should think [*S*] to intend that (1') be false. (Ibid.)

(It is clear from the context in which this appears that '*A*'s determination of *r*' means (roughly) the same as 'the inference (theoretical or practical) by which *A* reaches his response *r*'.)

I have my doubts as to whether this condition is necessary;⁴ more importantly, I fail to see how it eliminates all of the counter-examples it was designed to eliminate. In the "Moon Over Miami" example, for instance, it is true that *S* intends *A*'s belief that *S* wants *A* to leave to be an "inference-element" in *A*'s practical inference and true that *S* intends that *A* think that *S* does not intend *A*'s belief that *S* wants *A* to leave to be an "inference-element" in *A*'s practical inference; nevertheless, there is no reason to suppose that *A* thinks (or was intended to think) that *S* intends *A* *not* to have his belief that *S* wants him to leave as an "inference-element" in his determination to leave: this example would still be a counter-example if *A* thought (and was intended to think) that *S* sang "Moon Over Miami" not caring in the least if *A* left because *S* wanted him to leave but simply expecting that a bit of repulsive singing would be more effective.

I will not attempt to construct a more effective condition of this type, for I think that no condition requiring *S* not to have certain intentions will adequately deal with the problems raised by the preceding counter-examples. But before offering what I

⁴ It follows from what I say in section II. 3 that this condition cannot be a necessary condition.

think is a solution to those problems, I should like to consider one more alleged counter-example, one put forward by John Searle.⁵

(d) An American soldier (*S*) is captured during World War II by Italian troops (*A*). *S* would like to get *A* to believe he is a German officer by telling them in German or Italian that he is a German officer, but he knows neither language. However, *S* does remember a line from a German poem he memorized in school, and, hoping *A* does not know German, he utters the sentence, 'Kennst du das Land, wo die Zitronen blühen?' ('Knowest thou the land where the lemon trees bloom?') with the intention of deceiving *A* into thinking that he is saying in German "I am a German officer".

Thus, *S* utters the sentence 'Kennst . . .' intending

- (1) that 'Kennst . . .' be a German sentence (*f*);
- (2) that *A* recognize that 'Kennst . . .' is a German sentence;
- (3) that *A* infer at least in part from the fact that *S* uttered a German sentence that *S* uttered 'Kennst . . .' intending (4):
- (4) that *S*'s utterance of 'Kennst . . .' produce in *A* the belief that *S* is a German officer;
- (5) that *A*'s recognition of *S*'s intention (4) shall function as part of *A*'s reason for believing that *S* is a German officer;

and since there is no reason why *S* should not want *these* intentions recognized, or even recognized that they are intended to be recognized, we may consider him to have intentions (6) and (7) too. Nevertheless, I think one feels that it would be wrong, or at least not quite right, to say that *S* meant that he was a German officer (or anything else) by uttering 'Kennst . . .'

Searle's explanation of what precludes this case from being an instance of *S*-meaning is this:

We have here a case where I am trying to produce a certain effect by means of the recognition of my intention to produce that effect, but the device I use to produce this effect is one which is conventionally, by the rules governing the use of that device, used as a means of producing quite different illocutionary effects.⁶

⁵ "What is a Speech Act?", pp. 229–30; *Speech Acts*, p. 44.

⁶ "What is a Speech Act?", p. 230.

This explanation will not do. At a boring party Miss *S* might say to her escort, Mr. *A*, 'Don't you have to inspect the lemon trees early in the morning?', and mean thereby that she wants to leave.

What, if anything, makes Searle's example a counter-example? Strictly speaking, Searle's example is under-described, and it is possible that the American soldier uttered 'Kennst . . .' intending his Italian audience to reason (roughly) in either of the following two ways. (i) "'Kennst . . .' is a German sentence. *S* knows (or should know) that we do not speak German and knows that we know he knows this. But given the fact that he uttered a German sentence in these circumstances he is most likely uttering a sentence which means "I am a German officer" in German with the intention that we should infer that he intends to inform us that he is a German officer simply in virtue of the fact that he uttered a German sentence in these circumstances." However, I think the more likely interpretation of the soldier's scheme is that he intended his audience to reason: (ii) "*S* uttered 'Kennst . ..'. This is a German sentence. The fact that he uttered a German sentence shows that he thinks or hopes we understand German. So he must intend us to recognize, in virtue of our knowledge of German, what the sentence means and to infer from this what he meant by uttering it. . . . Well, given the preceding and the circumstances, it is most likely that 'Kennst . . .' means the same as 'I am a German officer' and that *S*'s intention in uttering this is to inform us that he is a German officer."

If we think of the American soldier as intending his captors to reason in the way characterized in (i), then I think it will not be incorrect to say that he did mean that he was a German officer. For in this event Searle's example will not differ relevantly from the following example of Grice's, where, I believe, we should want to say that something was meant.

A Port Said prostitute, lounging seductively in a shady doorway, utters in ingratiating tones an Arabic sentence which means "You filthy pig of a sailor" with the intention of picking up a British sailor.⁷

Unlike Searle's example, we are here inclined to say that the prostitute did mean something by uttering the Arabic sentence, and this suggests that the relevant difference between the two

⁷ A bowdlerized version of this example is in "Utterer's Meaning and Intentions", p. 162.

cases, and the difference which accounts for our differing intuitions about the two cases, is that the soldier intended his captors to reason in the way characterized in (ii).

There are at least two possibly relevant respects in which (ii) differs from (i).

(1) In both cases *S* utters x intending *A* to think *falsely* that x has a certain feature f' (that 'Kennst . . .' means "I am a German officer"), but it is only in (ii) that *S* intends *A* to have this false belief that x is f' as part of his reason for thinking that *S* uttered x intending to produce a response r in *A* (etc.).

But that this is *not* a sufficient condition for *S* not meaning something by uttering x is shown, I believe, by the following example.

S's wife, a present-day Mrs. Malaprop, confuses the words 'erotic' and 'erratic' so that she believes that 'erotic' means "erratic" and that 'erratic' means "erotic". *S* would like to tell his wife not to order his buttermilk from the milkman, because the milkman is too erratic. Rather than begin a futile explanation or use several words instead of one, *S* takes the easy way out and utters the sentence, 'Dear, please don't order my buttermilk from the milkman; he's too erotic'.

I do not think one would be at all disinclined to say that *S* here meant something and that what he meant by uttering 'he's too erotic' was that the milkman was too erratic, and this despite the fact that *A* was intended to have the false belief about the meaning of 'erotic' as her basis for inferring what *S* meant from what *S* uttered.

(2) I believe that what does make Searle's example a counter-example is this: *S* utters x intending *A* to think that x has a certain feature f (that 'Kennst . . .' is a German sentence) and intending *A* to think (wrongly)—on the basis of thinking x is f —that *S* uttered x intending *A* to think that x has a certain feature f' (that 'Kennst . . .' means "I am a German officer") and intending *A* to think—on the basis of *A*'s thinking that x is f' —that *S* uttered x intending to produce a response r in *A* (etc.). In other words, there is an intended discrepancy between *the basis A is intended to have for thinking S uttered x intending to produce r in A* and *the basis A is intended to think he is intended to have for thinking S uttered x intending to produce r in A*. (In the "Mrs. Malaprop" example there is no such discrepancy: the belief that 'erotic' means "erratic" is both the basis *A* is intended to have and the basis *A* is intended to think *A* is intended to have, and so on.)

It will be to no avail to revise Grice's definition by requiring S to intend A to recognize that S intends A to infer S 's intentions in uttering x from the fact that x has a certain feature(s) f ; for we could then construct a further counter-example in which there is an intended discrepancy between the basis A is intended to think he is intended to have and the basis A is intended to think he is intended to think he is intended to have. For despite first appearances, Searle's example is a counter-example of the same type as the preceding three.

II.2 Mutual knowledge*

I should now like to argue that there is a very common, ordinary feature of our everyday life, one which has to do with interpersonal knowledge, which once noticed will provide us with a condition which is at once a necessary condition for performing an act of communication and a condition which will eliminate those counter-examples based on deception without, I think, entailing a regressive series of intentions that "lower-order" intentions be recognized; although, as we will see, it does entail a quite harmless regress of the sort involved in knowing that one knows that p .

The phenomenon I am alluding to has no name, nor is there an otherwise simple way of referring to it. For this reason I shall coin the barbarism "mutual knowledge*", and I shall speak of two people, S and A , mutually knowing* that p , or of it being mutually known* that p by S and A , and so on; it will later be clear that any number of people greater than two may also mutually know* that p .

(a) Definition of 'S and A mutually know* that p'.

For convenience, let

' $K^*_{SA}p$ ' = df. 'S and A mutually know* that p'.

We can now say

$K^*_{SA}p$ iff
 $K_S p$ [S knows that p]
 $K_A p$
 $K_S K_A p$
 $K_A K_S p$
 $K_S K_A K_S p$
 $K_A K_S K_A p$

$K_S K_A K_S K_A p$
 $K_A K_S K_A K_S p$

·
 ·
 ·

(b) *Example.* Suppose that you and I are dining together and that we are seated across from one another and that on the table between us is a rather conspicuous candle. We would therefore be in a situation in which I am facing the candle and you, and you are facing the candle and me. (Consequently, a situation in which S is facing the candle and A , who is facing the candle and S , who is facing the candle and A , who is facing the candle and S , who is facing . . .) I submit that were this situation to be realized, you and I would mutually know* that there is a candle on the table.

Let us pretend that we are now in the situation described.

Clearly, I know that there is a candle on the table. So

$K_S p$.

I also know that you know that there is a candle on the table. How do I know this? First, I know that if a "normal" person (i.e., a person with normal sense faculties, intelligence, and experience) has his eyes open and his head facing an object of a certain size (etc.), then that person will see that an object of a certain sort is before him. Secondly, I know that you are a "normal" person and I see that your open-eyed head is facing the candle. (This is not intended to be overly exact.) So

$K_S K_A p$.

Further, I do not presume to be the only person aware of the above-mentioned law about normal people in certain circumstances; I also know that you know that normal people see things that are in their line of vision when their eyes are open, etc. And I have seen that you see that my open-eyed head is facing the candle. So I know that you know that I know that there is a candle on the table; i.e.,

$K_S K_A K_S p$.

Just as I know that you know the relevant law about normal observers, so I know you have just the same type of grounds for

knowing that I know it, too. And since I see that you see my head facing your head which has the candle in its line of vision, I know that you know that I know that you know that there is a candle on the table. So

$$K_S K_A K_S K_A p.$$

I could go on now to construct a further step which illustrates that for the same reason that I am justified in thinking that $K_A K_S K_A p$, so you have the same reason, *mutatis mutandis*, for thinking that $K_S K_A K_S p$, which would show, of course, that

$$K_S K_A K_S K_A K_S p;$$

but there is, I trust, no need to do this, for at this point it should be clear (i) that I can go on like this forever; (ii) that this regress is perfectly harmless; and (iii) that the phenomenon which obtains in this case is a general one: it will obtain, broadly speaking, whenever S and A know that p , know that each other knows that p , and all of the relevant facts are "out in the open".

(c) *Conditions for mutual knowledge**. Not quite so roughly, I think one can account for the possibility of mutual knowledge* and the conditions which must obtain for it to be realized in something like the following way.

First, I think that, in general, for any person x and any proposition p , if x knows that p , then there is a property H such that x knows that x is H and such that being H is sufficient for knowing that p , being H is sufficient for knowing that being H is sufficient for knowing that p , and so on;

$$\text{i.e., } (x) (p) (Kxp \rightarrow (\exists H) (KxHx \ \& \ (y) (Hy \rightarrow Kyp \ \& \ Ky(z) (Hz \rightarrow Kzp) \ \& \ Ky(z) (Hz \rightarrow Kz(w) (Hw \rightarrow Kwp)) \ \& \ Ky(z) (Hz \rightarrow Kz(w) (Hw \rightarrow Kw(v) (Hv \rightarrow Kvp))) \ \& \ \dots$$

For example, all "normal" people know that snow is white, know that all normal people know that snow is white, know that all normal people know that all normal people know that snow is white, and so on *ad infinitum*. (Likewise, I should think, for all or most of our common general knowledge; so if S and A mutually know* that each is "normal", all of the general knowledge each has in virtue of being a "normal" person will also be mutually known* by them.) To take another example, Jones knows that there is a house in front of him; he knows that anyone else in a

relevantly similar position will know that there is a house in front of him; he knows that anyone else in a relevantly similar position will know that anyone else in a relevantly similar position will know that there is a house in front of him, and so on.

This should not be surprising. In general, if one knows that p , one will know *how* one knows that p ; indeed, in most cases in which one knows how one knows that p one would not know that p unless one knew how one knew that p ; but knowing how one knows that p can be construed as just knowing that having a certain property is sufficient for anyone's knowing that p . Thus, when in such cases one knows that p , there is a property H such that one knows that being H is sufficient for knowing that p . But, in such cases, it will be in virtue of knowing that p that one knows that one knows that p because one is H , and so one will know that being H is sufficient for knowing that being H is sufficient for knowing that p ; and so on. Now, since relevant property H has this power with respect to knowing that p by virtue of the fact that being H is sufficient for knowing that p , the preceding line of argument suggests another truth about properties which satisfy the above conditions, a truth which will shortly be useful, viz.: there are properties, let H be one, such that, for any proposition p , if being H is sufficient for knowing that p , then being H is sufficient for knowing that being H is sufficient for knowing that p ; i.e., $(\exists H) (p) ((x) (Hx \rightarrow Kxp) \rightarrow (y) (Hy \rightarrow Ky(z) (Hz \rightarrow Kzp)))$.

Given the above, we can account for how it is that we mutually know* that there is a candle on the table. (No doubt we would mutually know* that each of us is "normal" before we got to the table, but for present purposes I shall pretend this is not the case.)

First, all "normal" people know that if a normal person has his eyes open and his head facing an object of a certain type (etc.), that person will know that an object of that type is before him; all normal people know that all normal people know this; and so on.

Second, all "normal" people know that anyone who appears and/or behaves in a certain way Ψ is normal; all normal people know that all normal people know this; and so on.

Now in the above example we have a situation in which, *ex hypothesi*, it is the case that

- (1) *S* is (visibly) ψ and facing the candle and *A*;
- (2) *A* is (visibly) ψ and facing the candle and *S*.

From this, together with the preceding two assumptions, it follows (once we fill in the obvious details) that *S* and *A* mutually know* that there is a candle on the table.

A situation in which two people are face-to-face with one another and the fact that *p* is a paradigm of mutual knowledge*, if anything is, but to labour, I hope, the obvious, it is not the only type of situation which generates mutual knowledge*. In particular, two people might mutually know* that *p* even though they are not directly acquainted with one another and even though they each have entirely different grounds for thinking that *p*. Thus, two people who have never met but who know of each other may reasonably assume that they mutually know* that London is a city in England; it might be mutually known* by *S* and *A* that *S* is in pain because it is mutual knowledge* that *S* is behaving in a certain way; or, to take an example closer to home, since *S* and *A* mutually know* that *S* uttered the sentence 'Please pass the salt', they mutually know* that *S* uttered 'Please pass the salt' intending that *A* should pass the salt because *S* wants him to.

In each case of mutual knowledge* we have a finitely describable situation such that in virtue of certain general features of this situation it follows that two people have an infinite amount of knowledge about each other. We ought, then, to be able to state a set of conditions which are such that *S* and *A* will mutually know* that *p* just in case these conditions are satisfied. The above discussion suggests that we might capture just those general conditions which generate mutual knowledge* in the following way. First, allow that it is a truth about knowledge in general that, for any property *H* and any proposition *q*, if one knows that whoever is *H* knows that *q*, and if one knows of any particular person that he is *H*, then one knows that that person knows that *q*; i.e., $(x)(H)(q)(z)((Kx(y)(Hy \rightarrow Kyq) \& KxHz) \rightarrow KxKzq)$. We may now say that

S and *A* mutually know* that *p* iff there are properties *F* and *G* such that

- (1) *S* is *F*;
- (2) *A* is *G*;
- (3) both being *F* and being *G* are sufficient for knowing

that *p*, that *S* is *F*, and that *A* is *G*; i.e., $(x)(Fx \vee Gx \rightarrow Kxp \& KxFS \& KxGA)$;

- (4) for any proposition *q*, if both being *F* and being *G* are sufficient for knowing that *q*, then both being *F* and being *G* are sufficient for knowing that both being *F* and being *G* are sufficient for knowing that *q*; i.e., $(q)((x)(Fx \vee Gx \rightarrow Kxq) \rightarrow (y)(Fy \vee Gy \rightarrow Ky(z)(Fz \rightarrow Kzq) \& Ky(w)(Gw \rightarrow Kwq)))$.

These conditions are fairly abstract. Here is how they might be instantiated in an actual case of mutual knowledge*, for instance our "candle" example. In that example, let *F* = the property of being an object *x* such that *x* is a visibly "normal", open-eyed, conscious person who is identical with *S* and who, at a close distance, is directly facing the candle and *A*, who is a visibly "normal", open-eyed, conscious person who is directly facing the candle and *S*; and let *G* = the property of being an object *x* such that *x* is a visibly "normal", open-eyed, conscious person who is identical with *A* and who, at a close distance, is directly facing the candle and *S*, who is a visibly "normal", open-eyed conscious person who is directly facing the candle and *A*. Now, *ex hypothesi*, *S* is *F* and *A* is *G*, and I submit that by virtue of their respectively possessing properties *F* and *G*, both *S* and *A* know that there is a candle on the table, that *S* is *F* and that *A* is *G*. With respect to the satisfaction of condition (4), I have already argued that generally when there is knowledge there is some property *H* such that, for any proposition *q*, if being *H* is sufficient for knowing that *q*, then being *H* is sufficient for knowing that being *H* is sufficient for knowing that *q*. I suggest that if any properties are of this sort, properties *F* and *G* are; and to get from here to the desired conclusion that condition (4) is satisfied, I suggest the following principle: if two properties entail just the same purely general properties, and if both properties are such that, for any proposition *p*, it is true of each property that if it is sufficient for knowing that *p*, then it is sufficient for knowing that it is sufficient for knowing that *p*, then, for any proposition *q*, if either property is sufficient for knowing that *q*, then the one is sufficient for knowing that the other is sufficient for knowing that *q*. Thus, we may conclude that, for any proposition *q*, if both being *F* and being *G* are sufficient for knowing that *q*, then both being *F* and being *G* are sufficient for knowing

that both being F and being G are sufficient for knowing that q ; for, roughly put, the two properties do not differ relevantly, since S and A share all of the purely general properties entailed by either F or G .

(d) *Prima facie objections.* I think one would be hard put to find reasons why in the "candle" example I could not go on indefinitely in the way specified. But even granting this there are two objections which might be made.

(1) "What follows from the example is not that S and A mutually know* that there is a candle on the table, but only that by proceeding in the above way each is capable of *acquiring* an indefinite amount of knowledge of the kind specified."

(2) "Even if S and A *now* mutually know* that there was a candle on the table, it doesn't follow that they knew this before it was *pointed out* to them that they could go on indefinitely in the way indicated."

I will not attempt to answer these objections conclusively; but I think I can remove a good part of the motivation behind them by calling attention to the following two facts. First, it is no objection to the claim that S knows that p that the thought that p never once entered S 's head. For example, I trust that it is true of each philosophy don in Oxford that he knows that his maternal grandmother was never married to Benito Mussolini. Second, it is no objection to the claim that S knows A knows S knows A knows S knows that p that S may have to be "convinced" or "brought to see" that he is entitled to claim to know this. I doubt that many a non-philosopher would agree right off the bat that he knows that he knows that he knows that $843 + 2 = 845$. Moreover, if having to be "brought to see" that p is sufficient for defeating the claim that S knows that p , then this will be embarrassing in general for a Gricean account of communication. For a good deal of arguing is often needed to show someone that his intention in uttering 'Shut the door!' was to get his audience to shut the door by means of recognition of this intention.

(e) *A revised analysis.* In Strawson's example it is not mutually known* by S and A that S uttered x , and so not mutually known* by S and A that S uttered x intending to produce r in A by means of recognition of intention. In the "Moon Over Miami" and "Tipperary" examples it is mutually known* by S and A that S uttered x but not mutually known* by them that S uttered

x intending to produce r in A by means of recognition of this intention. In Searle's example it is not mutual knowledge* that intends his utterance x to be evidence that he uttered x intending to produce r in A by means of recognition of intention simply in virtue of x being a German sentence uttered in the circumstances.

In the standard or paradigm case of Gricean communication it is mutually known* by S and A that S 's utterance x has a certain feature(s) f and mutually known* that the fact that S 's utterance x is f (together with certain other facts) is conclusive evidence that S uttered x intending to produce a response r in A by means of recognition of this intention and conclusive evidence that S uttered x intending it to be mutual knowledge* that S 's utterance x is conclusive evidence that S uttered x intending to produce r in A by means of recognition of this intention.

This is what should be expected just on the basis of the fact that by making sure everything is out in the open S increases his chances of securing his primary intention. But I believe that we can say something even stronger: given that S has no intention of deceiving A , S can, in general, utter x intending to produce a certain response r in A by means of A 's recognition of this intention *only if* S expects it to be mutually known* by S and A that S uttered x intending to produce r in A by means of recognition of this intention.

In section II. 1(b) I argued that given that S does not intend to deceive A and given that S utters x intending to produce a response r in A by means of A 's recognition of this intention, S can expect A to think that S uttered x intending to produce r in A *only if* S expects A to think S intends to produce r in A by means of recognition of this intention. Grant that this is so, and suppose that S utters x intending to produce r in A by means of recognition of intention and that S has no intention to deceive A . Given, then, that S uttered x with the complex Gricean intention, he must think that A will (or might) think that S uttered x intending to produce r in A by means of recognition of this intention. But A knows that a necessary condition of S uttering x intending to produce r in A by means of recognition of this intention is that S expects A to think that S uttered x intending to produce r in A by means of recognition of intention. So A will not think that S uttered x intending to produce r in A by means of recognition of intention unless A thinks that S thinks that A thinks that S uttered x intending to produce r in A by means of

recognition of intention. *S*, knowing this is so, must thus think that *A* will think that *S* thinks that *A* will think that *S* uttered x intending to produce r in *A* by means of recognition of intention. *A*, knowing this is so, will not think that *S* uttered x intending to produce r in *A* by means of recognition of intention unless he thinks *S* thinks he thinks *S* thinks he thinks *S* uttered x intending to produce r in *A* by means of recognition of intention; and so on.

It might seem, then, that the way to revise Grice's analysis is to require that *S* utter x intending it to be mutually known* by *S* and *A* that the fact that *S* uttered x in the circumstances is evidence that *S* uttered x intending to produce a response r in *A* by means of recognition of this intention. But were we to revise the definition in this way we would still not have adequately dealt with the type of counter-example we are concerned to eliminate; for we would be open to a counter-example in which *S* utters x intending it to be mutually known* by *S* and *A* that the fact that *S* uttered x in the circumstances is evidence that *S* uttered x intending to produce r in *A* by means of recognition of intention, but also intending *A* to think that *S* intended it not to be mutual knowledge* that his utterance of x in the circumstances was evidence that *S* uttered x intending to produce r in *A* by means of recognition of intention. For example, *S* might arrange bogus "evidence" that *p* knowing that *A* will disclose his presence to *S* while *S* is contriving the "evidence" and knowing that it will, in that event, be mutual knowledge* that the fact that *S* is contriving "evidence" that *p* is evidence that *S* is doing this with the same type of scheme he had when he was the subject of Strawson's counter-example (we may suppose, too, that *S* intends *A* to think that *p* on the basis of thinking *S* intends *A* to think that *p*, despite the disclosure); and so *S* and *A* will mutually know* (or believe*) that *S* intends to produce in *A* the belief that *p* by means of recognition of this intention on the basis of mutually knowing* (or believing*) that *S* was contriving bogus evidence that *p*; yet we should not want to say that *S* meant that *p* by contriving bogus evidence that *p*.

Nor will it do to revise Grice's definition by requiring *S* to intend it to be mutual knowledge* that the fact that *S* uttered x in the circumstances is evidence that he intended the fact that he uttered x in the circumstances to be mutually known* to be evidence that he uttered x intending to produce a response r in *A* by

means of recognition of intention. For this leaves open the possibility of a counter-example in which *S* utters x with those intentions, but also intending *A* to think that *S* intended it not to be mutual knowledge* that the fact that *S* uttered x in the circumstances is evidence that he intended the fact that he uttered x in the circumstances to be mutually known* to be evidence that he uttered x intending to produce r in *A* by means of recognition of intention.

Most generally described, the problem seems to be this. So long as a proposed account of *S*-meaning requires *S* to have some intention *i* such that the analysis does not secure that it is impossible for *S* to intend *A* to think that *S* did not utter x with intention *i*, then the analysis will not be sufficient. In other words, if *S* is to mean something by uttering x , then all of the intentions necessary for his meaning something must be out in the open; there must be no possibility of "hidden" intentions which are constitutive of an act of meaning something.⁸

I suggest that at this stage we revise Grice's account of *S*-meaning in the following way.

S meant something by (or in) uttering x iff *S* uttered x intending thereby to realize a certain state of affairs *E* which is (intended by *S* to be) such that the obtainment of *E* is sufficient for *S* and a certain audience *A* mutually knowing* (or believing*) that *E* obtains and that *E* is conclusive (very good or good) evidence that *S* uttered x intending

- (1) to produce a certain response r in *A*;
- (2) *A*'s recognition of *S*'s intention (1) to function as at least part of *A*'s reason for *A*'s response r ;
- (3) to realize *E*.

A few comments may be helpful.

(1) Typically, *E* will essentially involve the fact that *S*, a person having such-and-such properties, uttered a token of type *X* having a certain feature(s) *f* in the presence of *A*, a person having such-and-such properties, in certain circumstances *C*. For instance, *E* might be, in part, the state of affairs which obtained when *S* uttered 'Please pass the salt' in the direction of *A* while they were dining together. From the fact that this state of affairs obtained we could reasonably conclude that *S* and *A* mutually knew* that it obtained and mutually knew* that the fact that this state of

⁸ Cf. Strawson, "Intention and Convention in Speech Acts", p. 454.

affairs obtained was conclusive evidence that *S* uttered 'Please pass the salt' intending *A* to pass the salt because *S* wanted him to pass the salt and intending to realize the aforementioned state of affairs.

(2) It follows from the original Gricean account of what it is for *S* to mean something by uttering *x* that *S* meant something by uttering *x* only if *S* uttered *x* intending thereby to realize a certain state of affairs *E'* such that *A* would recognize that *E'* obtained and recognize (think) that *E'* is evidence that *S* uttered *x* intending to produce a response *r* in *A*. I have argued that *S* can non-deceptively utter *x* with the original complex Gricean intention only if *S* intends it to be mutually known* (or believed*) by *S* and *A* that *S* uttered *x* with the complex Gricean intention. A consequence of this is that *S* can non-deceptively utter *x* with the complex Gricean intention only if *S* utters *x* intending thereby to realize a certain state of affairs *E'* such that it will be mutually known* (or believed*) by *S* and *A* that *E'* obtains and mutually known* (or believed*) by them that *E'* is evidence that *S* uttered *x* intending to produce a certain response *r* in *A* by means of recognition of this intention. Now an argument of the same type as that used to show that, necessarily, if *S* utters *x* non-deceptively intending to produce a response *r* in *A* by means of recognition of intention, then *S* intends it to be mutually known* (or believed*) by *S* and *A* that *S* uttered *x* intending to produce *r* in *A* by means of recognition of intention can be given to show that *A* will not take *E'* as evidence that *S* intends to produce *r* in *A* by means of recognition of intention unless *A* thinks that *S* expected it to be mutually believed* by *S* and *A* that *S* uttered *x* intending to realize *E'* (for *A* will not, *ceteris paribus*, think *E'* is evidence that *S* uttered *x* intending to produce *r* in *A* (etc.) unless he thinks *S* intended to realize *E'*; so *A* knows that *S* will expect *A* to think that *S* intends to realize *E'*; so *A* knows that *S* will expect *A* to think that *S* will expect *A* to think that *S* intends to realize *E'* . . .). It follows, then, that either *E'* will be intended by *S* to be mutually believed* to be evidence that *S* uttered *x* intending to realize *E'* or else there will be some other state of affairs *E''* which *S* also intends to realize and which *S* intends to be mutually believed* to be evidence that *S* uttered *x* intending to realize *E'*. If the first disjunct is the case, then we have the proposed definition; if the second disjunct is the case, then we have the proposed definition or a vicious regress.

(3) An intuitive reaction to the counter-examples based on deception is that what precludes these cases from being instances of *S*-meaning is that *S* does not utter *x* expecting that if the intentions with which he uttered *x* are satisfied, *A* will recognize that *S* meant something by uttering *x*. One's next reaction is likely to be to disregard one's first intuition; for to say that *S* meant something by uttering *x* only if *S* intended or expected *A* to recognize that *S* meant something by uttering *x* seems a bit too circular to be either right or perspicuous. The proposed redefinition is such that *S* cannot utter *x* with the requisite intentions without expecting it to be mutually known* (or believed*) by him and *A* that he did so, if the intentions with which he uttered *x* are satisfied; that is, *S* cannot mean something by uttering *x* without expecting that if the intentions with which he uttered *x* are satisfied, *A* will recognize that he meant something by uttering *x*. That this is so is secured by the fact that for any state of affairs *E* satisfying the definition and for any proposition *q*, if the obtainment of *E* is sufficient for *S* and *A* mutually knowing* that *q*, then the obtainment of *E* is sufficient for *S* and *A* mutually knowing* that *S* uttered *x* believing that the obtainment of *E* would be sufficient for *S* and *A* mutually knowing* that *q*. So the proposed redefinition has the merit of confirming one's original intuition without being circular.

(4) I do not think that the mutual knowledge* condition commits *S* to having an infinite number of intentions, since one may intend *all* the consequences of one's act without intending *each* consequence. However, even if the definition did entail an infinitely regressive series of intentions, I do not think that this would be an objection to the definition; for the concept of mutual knowledge* would enable us to see why such a regress would be harmless and why such a regress would not preclude us from providing sufficient conditions for an instance of *S*-meaning.

(5) Does the definition allow *S* to utter *x* without actually intending to produce *r* in *A*? The notion of "good evidence" as it occurs in the definition is to be understood in its epistemic sense, and this entails that *S* cannot think that *E* is good evidence that he intends to produce *r* in *A* if he knows that he does not have this intention. Now I think that the sort of intention relevant to the definition is such that for any intention of this sort, *S* knows whether or not he has that intention. It follows that if *S*

does not have a certain intention, then he knows that he does not have that intention; and it follows from this that if *S* does not know that he does not have a certain intention, then he has that intention. We may conclude, then, that if *S* intends *E* to be good evidence that he intends to produce *r* in *A*, then he intends to produce *r* in *A*. However, should this line of argument ultimately prove unacceptable, we could easily circumvent the need for it by directly stipulating in the definition that *S* has those intentions which he intends *E* to be good evidence that he has.

Even if it is granted that the revisions made so far are both necessary and sufficient for dealing with those counter-examples based on deception, there are at least three different and further respects in which the most recently revised analysans (and, of course, Grice's original analysans) does not provide a set of jointly sufficient conditions for *S* meaning something by uttering *x*. The objections alluded to will be encountered and dealt with in the next chapter, when I formulate my own account of *S*-meaning.

II.3 *Some objections to the alleged necessity of Grice's analysans*

The central condition in Grice's analysis, the condition which gives the Gricean account of *S*-meaning much of its force and originality, is that to mean something is to intend to produce a certain response in an audience *by means of the audience's recognition of one's intention to produce that response*. In this section I hope to point out that there are standard instances of *S*-meaning where *S* does utter *x* intending to produce a certain response *r* in *A*, but where it is no part of *S*'s intention that part of *A*'s reason for his response *r* be that *S* intends to produce *r* in *A*. Pointing this out will show, in effect, that while the condition that *S* intend to produce *r* in *A* *by means of recognition of this intention* is perhaps a necessary condition for performing an act of "telling"—either "telling" *A* that such-and-such is the case or "telling" *A* to do such-and-such—it is not a necessary condition for *S* meaning something by uttering *x*.

(a) If Grice's account of what it is for someone to mean something were correct, an unwelcome and somewhat ironic consequence would be that although Grice will have written and published an article of several pages on what it is for someone to mean something, Grice will have meant almost nothing by what he wrote. For although Grice's primary intention in writing his

paper was to induce in us certain beliefs (albeit of a "conceptual" nature) about meaning, he neither expected nor intended that our reason for believing that what he wrote is true would be the fact that he intended us to believe that what he wrote is true. To put things somewhat crudely, we might say that in a philosophical-type argument *S* intends to get *A* to believe that *r* by pointing out that certain propositions, which *A* already believes, are reasons for believing *r*. But *S* does not intend *A* to think that *p* and *q* are reasons for believing *r* because *S* intends *A* to believe this; rather, *S* expects that *A* will see that what *S* said is true by virtue of *A*'s ability to reason from the premisses to the conclusion.

That there are such counter-examples to Grice's analysis is just what one should expect on the basis of certain *a priori* considerations. Grice's account of *S*-meaning is arrived at (in "Meaning") primarily by way of an abstraction from two types of cases: cases of "telling" someone that such-and-such is the case and cases of "telling" someone to do such-and-such. Thus, according to Grice, *S* meant that *p* by uttering *x* just in case *S* uttered *x* intending *A* to think that *p* on the basis (in part) of thinking that *S* intends *A* to think that *p*. It is implicit in Grice's account that *S* will intend *A* to take the fact that *S* intends *A* to think that *p* as evidence that *S* thinks that *p* (or, more commonly, thinks that he knows that *p*) and intend *A* to take the fact that *S* thinks (he knows) that *p* as evidence that *p* is the case. Suppose now that there are certain types of propositions such that it is well-known that, *ceteris paribus*, if *p*' is a proposition of this type, no one will take the fact that someone (no matter who) thinks he knows that *p*' as evidence that *p*' is the case. Since one cannot, in general, intend to bring about that which one knows one cannot bring about, it will follow from Grice's account of what it is for someone to mean that such-and-such is the case, that if *p*' is a proposition of this type, no one can mean that *p*'. But there is a large class of propositions for which ordinary empirical evidence (and so, *a fortiori*, the fact that a proposition of this type is believed by so-and-so) is irrelevant, at least in the sense that what is needed to believe that a proposition of this type is true is not new empirical evidence. This is true of most of the assertions made in journals of a theoretical or conceptual nature. For example, it is unlikely that anyone—ingratiating students apart—will believe that there are no analytical truths because he knows that Quine believes this to be true. So if Grice is right, no one could ever (or

hardly ever) mean that p' , where p' is a proposition of this type. But some people do say and mean that there are no analytical truths. So Grice is wrong. (We might express the objection being made by saying that Grice has overlooked the difference between sophists and cretans: if I put forward an argument intending to convince you that its conclusion is true, I cannot correctly be said to have *lied* to you, even though I know the conclusion is false and the argument invalid. To be a cretan I must intend you to think that p on my authority.)

This type of counter-example is not restricted to cases of *S* meaning that such-and-such is the case; it applies as well to cases of *S* meaning that *A* is to do such-and-such. For example, I could not *tell* you not to vote, but perhaps I could present you with a persuasive *practical* argument the conclusion of which is that you should not vote.

It is also relevant to note that a variation of the type of counter-example under discussion provides a counter-example to another claim of Grice's, viz., that if *S* utters x with the primary intention of producing in *A* the response r_1 by means of recognition of intention and also with the intention that the production in *A* of r_1 be sufficient for the production in *A* of some further response r_2 , then *S*'s intention to produce r_2 in *A* is no part of the characterization of what *S* meant by uttering x . In the cases considered in the preceding couple of paragraphs, *S* argues from premisses already believed by *A* to an unknown conclusion. But there are also cases where *S* tells *A* that p and intends *A*'s belief that p to be his reason for believing q (or doing ψ) and where *S* can nevertheless be said to have meant that q (or that *A* was to ψ). This might be the case were *S* to say, for example, "Don't swim in the water; it's shark infested", or "John's car broke down; so he won't be here".

(b) Counter-examples to the condition that *S* intend to produce r in *A* by means of recognition of intention are also provided by cases of *reminding* and *pointing out*.

- (i) *A*: "Now, what was that girl's name?"
S: Holds up a rose; or, "Rose".
- (ii) *A*: "All Balliol men are brilliant."
S: "Except Crumley-Gadswell."

In each of these cases *S* utters something x (utters 'Rose', or holds up a rose or utters 'Except Crumley-Gadswell') with the inten-

tion of thereby bringing it about that *A* is reminded that p ; yet it need be no part of *S*'s intention that the intended response be brought about by means of *A*'s recognition of *S*'s intention to produce that response; *S* knows that the mere utterance of x is enough to secure that *A* is reminded that p .

- (iii) *A*: "A necessary condition of someone's meaning that p is that he utter a sentence which means " p ."
S: "But then one could never mean that p by uttering a sentence metaphorically."

Here *S* utters x with the intention of getting *A* to see, or of pointing out to *A*, that p (that such-and-such is a consequence of *A*'s statement), but *S* does not, and indeed could not, intend to bring about this effect by means of recognition of intention. Still, we should want to say that *S* meant something.

In some of these cases *S* could be said to have the intention of producing in *A* the belief that p (e.g., where *A* had completely forgotten that p), but in other cases we should want to say that *A* already had (and was known by *S* to have) the latent belief that p , and in such cases it would be more accurate to say that *S*'s intention was to produce in *A* the *activated* or *occurrent* belief that p . Thus, in all cases so far considered of *S* meaning that p , *S* utters x intending to bring it about that *A* has in mind the belief that p , although only in some cases does *S* actually intend to provide *A* with the belief that p . Another type of case which fits this description would be a case in which *S* utters, e.g., 'Rose earns 30 pounds a week, has her own car and a flat in London' with the intention of reviewing or "calling up" these facts about Rose known to both *S* and *A* with the intention of securing that *A* has all of the relevant facts about Rose in mind while deciding whether to marry her.

Analogous examples may be thought of where the intended response is a practical one. For example, after their marriage, Rose might remind or point out to *A* that he should be doing the washing-up. (Perhaps in some cases of this type we should say, as Grice has suggested, that *S*'s intention is to restore *A* to activated intention.)⁹

I do not think that it can plausibly be denied that I have produced examples (i) in which *S* utters x with the primary intention of producing in *A* a certain response r and (ii) in which it is not

⁹ "Utterer's Meaning and Intentions", p. 171.

the case that *S* intends part of *A*'s reason for his response *r* to be based on *A*'s recognition of *S*'s intention to produce *r* in *A* and (iii) which would and should be classed as instances of *S*-meaning. This being so there seem to be only two lines of attack open to anyone who wishes to deny that I have produced genuine counter-examples to the condition requiring that the intended response be intended to be brought about by means of recognition of intention.

(1) In each example produced in this section *S* utters α intending to produce a certain response *r* in *A* but not intending to produce *r* in *A* by means of recognition of intention. It might be suggested that in one or more of the above types of cases there is some other response *r'* which *S* does intend to produce in *A* by means of recognition of intention and that it is in virtue of his uttering α with *this* intention that *S* means something by uttering α .

There may appear to be two candidate responses in the case where *S* argues that *p* from certain known premisses. (i) It has been suggested that while, in this type of case, *S* does not intend to get *A* to believe that *p* by means of recognition of intention, *S* does intend *A* to think that *S* thinks *p* and does intend that at least part of *A*'s reason for thinking that *S* thinks that *p* be that *S* intends *A* to think that *S* thinks that *p*. Perhaps, the suggestion goes, *S* means that *p* in virtue of his intention to produce in *A*, by means of recognition of intention, the belief that *S* thinks that *p*. (ii) When *S* presents an argument to convince *A* that *p* he certainly does not intend *A* to accept the argument on the basis of *S*'s accepting it, but he may intend *A*'s reason for *entertaining* the argument to be that *S* wants him to. So perhaps in this type of case *S* means something in virtue of his intention to get *A* to follow the argument, think of possible objections, etc.

It is even more difficult to find a candidate response in the reminding and pointing out cases, but here, too, it has been suggested that *S* intends *A* to think that *S* thinks that *p* and intends *A*'s thinking that *S* thinks that *p* to be based partly on *A*'s recognition of *S*'s intention to get *A* to think that *S* thinks that *p*.

We have, then, the merits of the following two suggestions to consider. (i) In the "argument" or "reminding" or "pointing out" cases *S* means that *p* by uttering α only if *S* intends to produce in *A* the belief that *S* thinks that *p* and intends at least part of *A*'s reason for thinking that *S* thinks that *p* to be that *S*

intends *A* to think that *S* thinks that *p*. (ii) In the "argument" cases *S* means that *p* by uttering α only if *S* intends *A* to entertain a certain argument the conclusion of which is that *p* and intends at least part of *A*'s reason for entertaining the argument to be that *S* intends *A* to entertain the argument.

I have two replies to these suggestions. First, even if these suggestions were true, this line of defence would constitute a significant retreat from Grice's original account. Any account of what it is for someone to mean something by uttering α must be inextricably tied to an account of *what* is meant by uttering α . Now it was essential to Grice's original theory that *S* meant that such-and-such was the case by uttering α just in case *S* uttered α intending to produce in *A*, by means of recognition of intention, the belief that such-and-such was the case. Thus, the above examples constitute at least counter-examples to Grice's account of what it is for someone to mean that such-and-such is the case, and, consequently, to adopt either of the above suggestions would amount to the adoption of a quite different theory. (This same objection can be made, *mutatis mutandis*, with regard to similar suggestions for an account of what it is for someone to mean that someone is to do such-and-such.) Secondly, both suggestions are palpably false. That (i) is false at least with regard to the "argument" cases may be seen by noting that in many instances where *S* argues that *p* it is already mutual knowledge* between *S* and *A* that *S* believes that *p*. For example, we both may know that I firmly and sincerely believe that this is the best of all possible worlds; my purpose in arguing that this is the best of all possible worlds is to convince *you* that it is. That (i) is false with regard to the "reminding" and "pointing out" cases may be seen by noting that in most cases in which *S* succeeds in reminding or pointing out to *A* that *p* he will also succeed in reminding or pointing out to *A* that *S* thinks that *p*. Suggestion (ii), on the other hand, would be most unattractive even if an intention to secure that *A* "entertains" a certain argument (or proposition?) were a constant feature of the type of case it is meant to accommodate; for it is clear that such an intention would only apply to a limited range of cases of *S* meaning that such-and-such is the case, and so taking this way out would commit one to saying that the conditions necessary for *S* meaning that such-and-such is the case vary from one type of case to the other. Moreover, since the "argument" and "pointing out" cases overlap, it would seem to

be false that whenever one argues that *p* from certain premisses one intends one's audience to "entertain" the argument put forward.

(2) A different line of attack would be to argue that the purported counter-examples are "derivative" from or "parasitic" upon cases of "telling" and are therefore to be explained in terms of and as departures from this "primary" case of *S*-meaning. I agree that there are cases which one should not count as counter-examples to Grice's account of *S*-meaning, even though in such cases *S* does not intend to produce a response in *A* by means of recognition of intention and even though one would be naturally inclined to say of such cases that *S* meant something. For example, one might say to a counter-suggestible child, "Do keep banging on your drum!" with the intention of getting him to *cease* banging on the drum. Here, I think, one would be inclined to say that one told the child to keep banging on the drum, and yet there is no response which one intends to bring about by means of recognition of intention: one intends the child to cease banging because he thinks one intends him to keep banging. I should think it wrong to consider this type of case a *counter-example* to Grice's theory; for it would seem that we are inclined to say that the speaker "told" the child to persist in banging the drum in virtue of his *pretending to tell* him this in the primary sense of the word 'tell'. But what kind of explanation could be given to show that the above examples are dependent upon cases of "telling" for their status as instances of *S*-meaning?

III

AN ACCOUNT OF *S*-MEANING

Introduction

WHAT is it for someone to mean something by (or in) uttering *x*? We began seeking an answer to this question by considering what seemed to be the most plausible account to date, that put forward by Grice. On Grice's view, to mean something by uttering *x* is just to utter *x* intending to produce in some person a certain type of response in a certain type of way. The only restraint on the type of response is that it must be something which is within the control of the audience, at least in the sense that it is the type of response for which the audience may have reasons. The only restraint on the way in which the response is to be produced by uttering *x* is that *S* must intend that at least part of *A*'s reason for his response *r* will be that *S* uttered *x* intending to produce response *r* in *A*. One knows what *S* meant if, and only if, one knows what response *S* intended to produce in *A*.

I want to retain the condition that *S* meant something by uttering *x* only if *S* uttered *x* intending to produce some response in *A*. On the assumption that this condition is necessary and on the assumption that the mutual knowledge* conditions adequately deal with the problems they were meant to deal with, this leaves at least two main problems to be solved before we shall have arrived at an adequate and correct account of *S*-meaning.

One problem, which has to do with the specification of what *S* meant, was noticed in section I. 3, where it was remarked that even if correct, Grice's suggestion that what is meant is determined by, and only by, the value of 'r' does not provide a criterion for determining what *S* meant, but only a criterion for determining what must be determined if one is to determine what *S* meant. No account of *S*-meaning will be complete unless it provides a definite means for specifying what *S* meant by uttering *x*.

A more fundamental problem arises out of the fact that the restraint placed by Grice on the way in which *S* must intend to