

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?

Schiffer

Meaning, Chapter 3

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

produce the relevant response r in A is much too restrictive, leaving out all but one type of instance of S -meaning. So we want to find a new way of completing the sentence form ' S meant something by uttering x only if S uttered x intending to produce in A a certain response r by means . . .'.

My procedure in this chapter is, first, to provide a definiens for each of the sentence forms 'by uttering x S meant that such-and-such was the case' and 'by uttering x S meant that A was to do such-and-such' (the bulk of this is done in sections III. 1 and III. 2). In section III. 3, I argue that S meant something by (or in) uttering x if and only if by (or in) uttering x S meant that such-and-such was the case or I, for some A , S meant that A was to do such-and-such. This is my solution to the first of the above-mentioned problems. Roughly, essentially, and apart from mutual knowledge* conditions, I argue that S meant that such-and-such was the case if and only if S uttered x intending to cause in A the activated belief that such-and-such was the case and intending his utterance x to cause A to have the activated belief that such-and-such was the case at least partly by virtue of A 's recognition (belief) that x is related in a certain way R to (the type of response) thinking that such-and-such was the case. In the definiens for 'by uttering x S meant that A was to do such-and-such' the intended response is A 's doing such-and-such; otherwise the two definitions are essentially identical. Hence, my solution to the second of the above-mentioned problems is that S meant something by uttering x only if S uttered x intending to produce in A a certain response r by means of A 's recognition that x is related in a certain way R to the type of response to which r belongs.

It will be convenient if, at the outset, the expressions ' S meant that such-and-such was the case' and ' S meant that A was to do such-and-such' are replaced, respectively, by the expressions ' S meant that p ' and ' S meant that A was to ψ '. The only allowable substituends for ' p ' will be expressions which could complete the sentence form 'John believes that . . .'; the only allowable substituends for ' ψ ' will be expressions which could complete sentences of the form 'John will . . .' or 'John intends to . . .'.

III.1 S meant that p

The examples so far considered in which S means that p divide up, broadly speaking, into two types of cases. On the one

hand we have cases of S "telling" A that p , where S utters x intending to produce in A the belief that p by means of A 's recognition of this intention, and, on the other hand, we have cases of reminding, pointing out, arguing from known premisses to an unknown conclusion, etc., where S utters x intending to produce in A the belief that p or the *activated* belief that p (i.e., intends A to have in mind the belief that p), but not intending to produce this response by means of recognition of intention. There is at least one thing all of these cases have in common, viz.: in each case S utters x intending thereby to cause in A the *activated* belief that p .

I propose to accept provisionally the condition, inherited in part from Grice, that S meant that p by uttering x only if S uttered x intending thereby to cause in A the activated belief that p . Its acceptance will cease to be provisional when, in sections III.3 and III.4, the alleged objections to this condition are dealt with.

Needless to say, this condition is not a sufficient condition for S meaning that p , if for no other reason than that it is met by each counter-example described in section II.1. But there still remains the question whether the mutual knowledge* conditions are merely required for S to be *telling* A that p or whether these conditions are required for all instances of S meaning that p . That the latter is the case is strongly suggested by what follows. We know that at least some instances of S uttering x with the intention of reminding A that p are instances of S meaning that p by uttering x . Now suppose it were claimed that a sufficient condition of S having meant that p by uttering x is that S uttered x intending to remind A that p . Each of the following examples would be a counter-example to that claim.

(i) S knows that A is trying to remember so-and-so's name. S , unbeknown to A , places a bowl of roses in A 's room, intending that A should see the roses and be reminded that so-and-so's name is Rose. A is intended to think that the roses were placed in his room purely for decorative reasons.

(ii) Again S places a bowl of roses in A 's room with the intention of reminding A that so-and-so's name is Rose. This time, however, S intends A to recognize that S placed the roses in A 's room intending to remind A that so-and-so's name is Rose, but also intending A to *think* that S intends A to think that S placed the roses in the room purely for decorative reasons.

(iii) Again S intends the roses to remind A that so-and-so's name is Rose, only this time S intends A to reason: " S intends me to think that he intends me to think that he put the roses in my room purely for decorative reasons, but I recognize that he really intends me to recognize that his intention in putting the roses in my room was to remind me that so-and-so's name is Rose."

In none of these examples is it the case that S meant that so-and-so's name was Rose. Clearly, similar examples may be constructed where S intends to point out to A that p or to convince A that p on the basis of A 's prior acceptance of certain premisses, or, for that matter, any type of case in which S utters x intending to produce in A the belief or activated belief that p . Let us begin, then, with the following set of necessary conditions.

(A) S meant that p by (or in) uttering x only if 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 cause in A the activated belief that p ;
- (2) to realize E .

The next example shows that the conditions of (A) are not jointly sufficient for S meaning that p .

S , a rather advanced neuro-physiologist, knows that by striking a certain chord on a piano he will emit a sound of a certain frequency which will set off a certain neuro-physiological process in the brain of any person of a certain type which will result in that person's remembering what word he first learned as a child. A , let us suppose, is S 's assistant; so the above—along with the fact that A is a person of the relevant type and the fact that S knows on independent grounds what word A first learned—is mutual knowledge* between S and A . Finally, suppose that S strikes the relevant chord intending to cause A to remember that '*Gesundheit*' was the first word A learned and intending the satisfaction of the mutual knowledge* conditions.

I do not believe that it would be correct to say that by striking the chord S meant that '*Gesundheit*' was the first word A learned. In order to see what precludes the neuro-physiologist from

having meant that '*Gesundheit*' was the first word A learned it will be helpful to contrast the preceding example with the following two examples.

(i) A is trying to remember what musical instrument Liberace plays. S strikes a piano chord with the intention of reminding A that Liberace plays the piano.

(ii) This is the same as (i), only instead of striking a piano chord S utters the sentence 'Liberace plays the piano'.

I believe that with regard to examples (i) and (ii) one is inclined to say that S meant that Liberace plays the piano: in (i) S meant this by (or in) striking the piano chord and in (ii) S meant this by (or in) uttering 'Liberace plays the piano'. What precludes S from having meant that p in the 'neuro-physiologist' example is, I believe, the way in which striking the piano chord is intended to cause A to remember that p . Specifically, the relevant difference between the "neuro-physiologist" example and the first "Liberace" example seems to be this. In both examples striking the chord is intended to cause A to remember that p but in the "Liberace" example, unlike the "neuro-physiologist" example, striking the piano chord is intended to cause A to remember that p as a result of (or by virtue of) A recognizing that there is a certain relation (association, connection, or correlation) between S 's striking the chord and the belief that p (i.e., the belief that Liberace plays the piano). This is also the relevant difference between the "neuro-physiologist" example and the second "Liberace" example: in (ii) S intends his utterance of 'Liberace plays the piano' to cause A to remember that Liberace plays the piano by virtue of A 's recognition that 'Liberace plays the piano' is a conventional means for making known one's intention to produce in an audience the belief (or activated belief) that Liberace plays the piano. Thus, the relevant difference between (i) and (ii) is that in (i) S intends his utterance x to be a cause of A 's believing p by virtue of A 's awareness of a certain "natural" association or relation between x and (the response-type) believing p and in (ii) S intends his utterance x to be causally efficacious by virtue of A 's awareness that x is conventionally correlated with the belief that p (or, to quiet Quinean qualms, correlated with (the response-type) believing p).

So it seems reasonable to think that an instance of S uttering x with the intention of thereby reminding A that p is an instance of S meaning that p by uttering x only if S utters x intending to

cause in A the activated belief that p and intending that his utterance x be a cause of A 's having the activated belief that p at least partly by virtue of A 's recognition (belief) that x is related in a certain way R to the belief that p (or, what for me will be the same, (the response-type) believing (thinking) that p). And if this is reasonable, then it will be reasonable to think that S meant that p by uttering x only if S uttered x intending to cause in A the activated belief that p and intending that his utterance x be a cause of A 's having the activated belief that p at least partly by virtue of A 's recognition (belief) that x is related in a certain way R to the belief that p . For the same type of considerations as those which apply to cases of "reminding" apply as well to the "pointing out", "calling up" ("fact review"), and "argument" cases; and in cases of "telling" A that p it is clear that there is at least one relation between x and the belief that p which S must intend A to recognize if S is to succeed in producing in A the belief that p , viz., that S uttered x intending to produce in A the belief that p . The conditions of (A) may therefore be replaced by those of (B).

(B) S meant that p by (or in) uttering x only if 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 cause in A the activated belief that p ;
- (2) satisfaction of (1) to be achieved, at least in part, by virtue of A 's belief that x is related in a certain way R to the belief that p ;
- (3) to realize E .

In Chapter I we saw the need, in restating Grice's definition, to make explicit the condition that S intend A to think that S intends to produce a certain response r in A on the basis of thinking that S uttered x . Since an object or event will be evidence that such-and-such is the case only *qua* being an object or event of a certain *sort*, the restatement of Grice's account made explicit reference to a special "feature(s) f ". Subsequently this explicit reference was submerged when the mutual knowledge* conditions were added, but only submerged and not erased: S 's utterance x

will be the key constituent of the state of affairs E and it will have this role in virtue of having certain of its features and not in virtue of other of its features. There are now two further points to be made apropos of the most recently added demand that S intend a certain relation to obtain between x and the intended response. (1) The fact that x is f will not serve as a means for making known to A the particular response S intends to produce in A unless the fact that x is f relates (or is thought by A to relate) x in a certain way to the type of response to which r belongs. (2) If, as claimed, condition (2) of (B) is a necessary condition of S meaning that p , then A will not take the fact that S uttered x as evidence that S meant thereby that p unless A thinks that x is related in way R to the belief that p . So the requirement that S intend x to be (thought by A to be) related in a certain way R to the belief that p will have, as it were, a double life: if S meant that p by uttering x and if S succeeded in satisfying the intentions with which he uttered x , then x will have been (thought by A to be) related in a certain way R to the belief that p and at least partly in virtue of this fact S 's utterance x will have been both a cause of A 's having the activated belief that p and evidence that S uttered x intending to produce in A the activated belief that p .

(B) still does not keep out quite as much as one might wish.

(1) There are cases where, by uttering x , S both means something and intends to cause A to believe that p , but where S does not mean that p . For example, we may not wish to say that in prescribing to A that he ought to believe that his wife was faithful, S meant that A 's wife was faithful, a proposition S may not wish to commit himself to. And, I suppose, in suitable circumstances one person might even command another person to believe, or to make it the case that he believes, that a certain proposition is true. (B) would commit us to saying that the person would mean that that proposition is true. (These examples are also counter-examples to the alleged sufficiency of Grice's analyses. This is one of the objections alluded to in Chapter II, p. 42.)

(2) Imagine a person who suffers from the complaint that whenever anyone utters the sound 'Urf!' he immediately becomes obsessed with the wholly irrational belief that his dog Rover is about to be run over. Now suppose that S takes delight in putting

poor A in this horrible state by uttering 'Urf!'. Even if all of the conditions of (B) are satisfied, I still think it would be incorrect to say that by uttering 'Urf!' S meant that Rover was about to be run over.

(3) (B) also commits us to saying that in presenting Salome with the head of St. John the Baptist on a charger, Herod meant that St. John the Baptist was dead.

Let us begin with (3). What are we to say about examples of this type? Well, one thing that might be said is that in presenting Salome with the head of St. John the Baptist, Herod meant that St. John the Baptist was dead. This does not strike me as a wildly implausible thing to say. Consider an analogous case.

(3a) A: "Let's play squash."

S: Holds up his bandaged leg.

Here, I think, one would say, intuitively, that by holding up his leg S meant that he could not play, or that he could not play because his leg was injured; yet it would seem that the only difference between (3) and (3a) which is possibly relevant is that the "inference" A has to make in the "bandaged leg" example is slightly less direct than in the case of St. John the Baptist's head, although in both cases one could make the relevant inference without any assistance on the part of S .

Grice has objected that while we may say that (in (3a)) S meant that he could not play squash by holding up his bandaged leg, he could *not* mean thereby that his leg was bandaged.¹ But, in the first place, even this is not an objection to the point I am trying to make, which is that there is no relevant difference between (3) and (3a), so that if we may say that S meant that he could not play squash, then—by parity of reason—we may say that Herod meant that St. John the Baptist was dead (it was not suggested that Herod meant that there was a severed head on his charger). In the second place, I think that it is false that S could not mean that his leg was bandaged by holding up his bandaged leg. Consider (3b).

(3b) A: "I've heard that your leg is bandaged. Is it true?"

S: Holds up his bandaged leg.

Here, I think, one would say that S meant that his leg was bandaged.

¹ "Utterer's Meaning and Intentions", p. 170.

However, if it should turn out to be desirable to eliminate examples of type (3), this would not be difficult to do. This can be done by specifying that the relation R between S 's utterance x and the belief that p which A is intended to recognize must not be such that S 's utterance x will provide A with *evidence* that p without the mediation of an intention on the part of S to produce in A , by uttering x , the belief that p . And if the reader finds it implausible to allow that Herod meant that St. John the Baptist was dead, he may treat the proposed way in which cases of this type may be excluded from the definition as an optional condition which he may add to the analysis should his intuitions so compel him. I shall not.

A few general remarks will place us in a position to deal with cases of types (1) and (2). If S meant that p by uttering x , then not only will S have intended to produce in A the belief (or activated belief) that p , S will also have intended that A have *reasons* for his belief that p (although only in some cases does S intend to *supply* A with reasons for believing that p). Now it is almost always the case that if one holds p as a reason for believing that p , one will think that p provides, in some sense, grounds for thinking that p is true. Let us call reasons which are held in this way "truth-supporting" reasons. One's reasons for believing p need *not* be truth-supporting; one might have moral or prudential grounds for thinking that p . This distinction gives us a means for eliminating counter-examples of type (1); for we may say that a necessary condition for S meaning that p by uttering x is that S intend A to have truth-supporting reasons for his belief that p . And the requirement that S intend A to have reasons for thinking that p will also serve to eliminate counter-examples of type (2).

It will be convenient at this point to introduce the following notational device. Let

' S intends to produce in A the response r/p ' = df.

' S intends to produce in A the response r for which he intends A to have the reason(s) p .'

If A is intended to have truth-supporting reasons for his response (i.e., his belief), then this will be marked by writing

' $p(t)$ '.

It may then be said that S means that p by uttering x only if

S intends that there be some ρ such that his utterance of x causes *A* to have the activated belief that $p/\rho(t)$.

The intuitive justification for this condition is at least twofold. First, it captures the demand that meaning and communication be *rational* in a certain way; secondly, it reflects the fact that communication, in general, aims at the production of knowledge and not merely belief. Consequently it may seem that a more intuitive formulation of the needed condition would be both possible and desirable. However, I feel that the above formulation of the needed condition will prove most useful in the end, partly because of two forthcoming considerations. One is that we shall find that a similar restriction is required in the definiens for 'by uttering x *S* meant that *A* was to ψ '; the other is that, in the general account of speech acts to be offered in Chapter IV, an essential ingredient in the determination of the "illocutionary force" of an utterance is the reason or reasons *S* intends *A* to have for the primary response *S* intends to produce in *A*, and so it is desirable to have a formulation which brings this feature explicitly into view.

Definition (B) may be revised accordingly to give us (C).

(C) *S* meant that p by (or in) uttering x *only if* *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) there to be some ρ such that *S*'s utterance of x causes in *A* the activated belief that $p/\rho(t)$;
- (2) satisfaction of (1) to be achieved, at least in part, by virtue of *A*'s belief that x is related in a certain way *R* to the belief that p ;
- (3) to realize *E*.

III.2 *S* meant that p and *S* meant that *A* was to ψ

The account of what it is for someone to mean that p so far arrived at requires at least one further minor adjustment before sufficiency may be claimed for its conditions; but before any further adjustments are made it will be well to offer the following

parallel account, also not quite sufficient, of what it is for someone to mean that *A* is to ψ .

S meant that *A* was to ψ by (or in) uttering x *only if* *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* mutually knowing* (or believing*) that *E* obtains and that *E* is conclusive (very good or good) evidence that *S* uttered x intending

- (1) there to be some ρ such that *S*'s utterance of x causes *A* to ψ/ρ ;
- (2) satisfaction of (1) to be achieved, at least in part, by virtue of *A*'s belief that x is related in a certain way *R* to (the act-type) ψ -ing;
- (3) to realize *E*.

There are a few features of this set of conditions which call for special comment.

(1) The demand that *S* intend *A* to have a reason for ψ -ing is the demand that the intended act be rational, and it is needed to rule out the same type of case as its counterpart in the definition of '*S* meant that p '. Thus one might imagine a man so neurotically terrified of cats that whenever he hears the sound 'meow' he begins running around in circles—he just can't help himself. Condition (1) above precludes us from having to say that by uttering 'meow' *S* meant that *A* was to run around in circles.

(2) The conditions for *S* meaning that p allow that *S* may mean that p by uttering x even though *S* intends his utterance x to provide *A* with reason(s) for thinking that p without the mediation of *S*'s intentions (e.g. the "Herod" example). Analogously, the conditions for *S* meaning that *A* is to ψ allow that *S* may mean that *A* is to ψ by uttering x even though *S* intends his utterance x to provide *A* with a reason for ψ -ing without the mediation of *S*'s intentions. Thus, by posting a vicious dog at the entrance to his home, *S* may have meant that *A* was to stay away. Once again, my intuitions do not balk at this result, and once again there is an easy way of eliminating such cases should it be desirable to do so: it may be done by specifying that the relation *R* between *S*'s utterance x and (the act-type) ψ -ing which *A* is intended to recognize must not be such that *S*'s utterance

x will provide A with a reason for ψ -ing without the mediation of an intention on the part of S to cause A to ψ .

(3) A disanalogy between the definitions of ' S meant that p ' and ' S meant that A was to ψ ' is that in the former but not the latter the intended response is an activated propositional attitude. Grice has suggested achieving symmetry by replacing A 's ψ -ing with an (activated) intention on the part of A to ψ as the intended response in the account of imperative utterances.² But while this would perhaps add a certain attractiveness to the account of S -meaning, it would also, I fear, open the way for certain counter-examples in which S intends to produce in A the intention to ψ but does not (and is known by A not to) intend A actually to ψ .

(4) Finally, a word about the definiendum, ' S meant that A was to ψ '. In ordinary discourse, to say that S meant that A was to ψ often implies something fairly strong, e.g. that S was ordering A to ψ . If I politely request you to pass the salt, it may be misleading to say that I meant that you were to pass me the salt. There is, I suppose, a certain amount of regimentation being imposed in our choice of a definiendum, but it should be noted that it is compatible with the account of what it is for someone to mean that A is to ψ that the value of ' ψ ' may be, say, passing the salt *if one pleases*.

As I remarked earlier, there is a further minor revision to be made before sufficiency may be claimed for the conditions of either definition. The importance of the adjustment to be made lies mainly in certain of its consequences, not all of which will be explored in this chapter.

Consider the most recent analyses for the account of what it is for S to mean that p by uttering x , i.e. (C) on p. 58. The problem is this. If the conditions of (C) were jointly sufficient for S meaning that p by uttering x and if S uttered x with the intentions specified in (C), then not only will S have meant thereby that p , he will also have meant by uttering x that he uttered x intending to cause in A the activated belief that p , that he intends A to have reasons for his belief that p , that he intends his utterance x to cause in A the activated belief that p at least partly by virtue of A 's recognition that x is related in a certain way R to the belief that p , and so on for all the other beliefs which (C) requires S to intend, in uttering x , to produce in A in addition to the belief (or activated belief) that p . But

² "Utterer's Meaning and Intentions", p. 166.

if S means, say, that the cat is on the mat by uttering 'The cat is on the mat', it will not be the case that in uttering 'The cat is on the mat' S will also have meant that he uttered 'The cat is on the mat' intending to produce in A the belief that the cat was on the mat. So the conditions of (C) are not jointly sufficient for S meaning that p by uttering x . The same type of objection can be made against the above account of what it is for S to mean that A is to ψ .

To see that the definitions have this consequence, suppose that S utters x with the intentions specified in (C). S will, in that event, intend to produce a certain state of affairs E which is mutually known* by S and A to be evidence that S uttered x intending to produce in A the belief (or activated belief) that p . Hence, S intends by uttering x to cause in A the belief that S uttered x intending to cause in A the activated belief that p , and since S intends his utterance x to be evidence that S uttered x intending to cause in A the activated belief that p , S intends his utterance x to be a cause of A 's having the belief that S uttered x intending to cause in A the activated belief that p by virtue of A 's recognition that x is related in a certain way to the belief that S uttered x intending to cause in A the activated belief that p . And since all of this is made known by the state of affairs E (viz., the fact that S uttered x in the circumstances), it follows that if (C) provided sufficient conditions, S will have meant by uttering x that he uttered x intending to cause in A the activated belief that p . Clearly, this is unacceptable.

It is worth noticing that Grice's account of S -meaning is open to the same type of objection. Suppose S utters x intending to inform A that p by means of recognition of this intention. In order to inform A that p , S will, in uttering x , intend to produce in A the belief that S believes that he knows that p ; but S will intend A 's reason for thinking that S thinks that he knows that p to be, at least in part, the fact that S uttered x intending A to think that S thinks he knows that p . So it is a *de facto* consequence of Grice's account that whenever S utters x intending to produce a certain response r in A by means of recognition of intention, there will always be certain beliefs, other than the primary response aimed at, which S also intends to produce in A by means of recognition of intention. So it follows from Grice's account that whenever S means something he will mean a lot more than he bargained for. This is the second respect alluded

to in section II.2, p. 42, in which Grice's analysis is not sufficient.

What we want to say is that if *S* means that *p* by uttering α , then, while he must intend to produce in *A* certain beliefs in addition to the belief (or activated belief) that *p*, these other intentions to produce beliefs do not constitute separate acts of meaning; rather they are necessary constituents of *S*'s act of meaning that *p*, or a necessary part of the characterization of *S*'s act of meaning that *p*. The problem is to find the relevant difference between *S*'s intention to produce in *A* the belief that *p* and, say, his intention to produce in *A* the belief that *S* intends to produce in *A* the belief that *p*. In other words: what is the relevant difference between, say, *S*'s intention to produce in *A* the belief that *p* and his intention to produce in *A* the belief that he intends to produce in *A* the belief that *p* such that in virtue of this difference *S* means that *p* by uttering α but does not mean by uttering α that he intends to produce in *A* the belief that *p*?

One relevant difference which comes quickly to mind (a difference of which I have been tacitly availing myself all along) is that when it is true to say that *S* meant that *p* by uttering α , then *S*'s intention to produce in *A* the activated belief that *p* is the *primary* intention with which *S* uttered α , whereas *S*'s intention to produce in *A* the belief that *S* uttered α intending to produce in *A* the activated belief that *p* is merely a *secondary* intention *S* had in uttering α —one carried on the back of the primary intention—an intention *S* had in uttering α as a result of having the primary intention to produce in *A* the activated belief that *p*.

Generally, when a person does an act *X* a distinction can be made between: (i) the intention(s) *with which* that person did *X*, and (ii) certain other intentions that person merely had *in* doing *X*. If an intention *i* is an intention with which one did *X*, then *i* will be an intention one had in doing *X*, but the converse need not hold. To specify one's primary intention in doing *X*—the intention with which one did *X*—is to give one's reason for doing *X*. This does not apply to all of the intentions one had in doing *X*.

I believe that this asymmetry applies to the difference between *S*'s intention to produce in *A* the belief (or activated belief) that *p* and the other beliefs *S* must intend to produce in *A* when he means that *p* (or that *A* is to ψ), and I believe it is this asymmetry

which accounts for why, in uttering α , *S* meant that *p* and not, say, that he uttered α intending to produce in *A* the belief that *p*. (*S* may, of course, utter α with more than one primary intention, and in such cases he may mean more than one thing; in some cases it may not be clear what *S*'s primary intention was, and to the extent that this is so it will not be clear what *S* meant.)

Finally, I submit the following two definitions.

S meant that *p* by (or in) uttering α iff *S* uttered α 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 α with the primary intention

(1) that there be some ρ such that *S*'s utterance of α causes in *A* the activated belief that $p/\rho(t)$;

and intending

(2) satisfaction of (1) to be achieved, at least in part, by virtue of *A*'s belief that α is related in a certain way *R* to the belief that *p*;

(3) to realize *E*.

S meant that *A* was to ψ by (or in) uttering α iff *S* uttered α 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* mutually knowing* (or believing*) that *E* obtains and that *E* is conclusive (very good or good) evidence that *S* uttered α with the primary intention

(1) that there be some ρ such that *S*'s utterance of α causes *A* to ψ/ρ ;

and intending

(2) satisfaction of (1) to be achieved, at least in part, by virtue of *A*'s belief that α is related in a certain way *R* to (the act-type) ψ -ing;

(3) to realize *E*.

The mutual knowledge* conditions will have to be adjusted slightly to accommodate certain utterances in the absence of an audience (this is done in section III.4), but, for all practical purposes, these definitions may be regarded as final, which is not, however, to say that they will encounter no further problems.

In fact, one such problem will now be raised and I hope that its solution will serve to confirm at least certain features of our definitions.

A consequence of the latest revision is that if S utters x with just those intentions specified in either definition, then he *cannot* mean by uttering x that his primary intention in uttering x is to produce in A the belief that p (or to get A to ψ). In view of the preceding, this may hardly seem objectionable. But suppose that S uttered x with just those intentions specified in the account of what it is for S to mean that p and suppose that what S uttered (the value of ' x ') was

σ_1 : My primary intention in uttering this sentence is to produce in you—by means of recognition of intention—the belief that p .

While it may not seem objectionable to say that in uttering σ_1 , S meant that p , it does *seem* objectionable to say that S did not also mean just what he said, viz., that his primary intention was . . . After all, we know what σ_1 means and we seem to understand what S would mean by uttering it; indeed, unless S is lying it would seem that what he said is true. But if S uttered σ_1 with just those intentions specified in the definiens for ' S meant that p ' and if by uttering σ_1 S meant that his primary intention was . . ., then we have a counter-example to the above definitions, a counter-example, as we will shortly see, with quite serious consequences for the entire programme we are engaged in.

It is essential to the programme of providing an account of the meaning of utterance-types in terms of a basic account of S -meaning that what S means by uttering x is not at all determined by what is uttered, i.e. by the value of ' x '. As I remarked in section I.3, this does not mean that one can do or say whatever one likes and mean thereby anything one pleases to mean: one must utter x with the relevant intentions, and not any value of ' x ' will be appropriate to this end. But whatever the value of ' x ' is, it is essential to our account of meaning that if one can reasonably expect to utter x with those intentions purported to be necessary and sufficient for meaning that p , then one will mean thereby that p , no matter what x is. Another way of making this point is this. If the only difference between two utterances is that in one case S utters x and in the other S utters y , then what S means by uttering x is identical with what

S means by uttering y . The importance of this condition is that if it were the case that what S meant by uttering x were determined, even in part, by the meaning of x , then this would, on the face of it, render circular an account of what x means in terms of what is or would be meant by uttering x . And the serious problem with the objection raised is that, if true, it would show that what is meant by uttering x at least sometimes is determined by the meaning of x .

For suppose that instead of uttering σ_1 , S had simply uttered

σ_2 : p [e.g. 'The cat is on the mat.']

And let us suppose that the only difference between the two utterances is in what S uttered. For instance, we might imagine that S desired to inform A , by means of recognition of intention, that the cat was on the mat. He therefore had to utter *something* which would make known to A his intentions. In the one case he utters 'My primary intention in uttering this sentence is to inform you—by means of recognition of intention—that the cat is on the mat' (σ_1), and in the other case he simply utters 'The cat is on the mat' (σ_2); but the only difference between the two cases is in the means used to make known his intentions, i.e. in the value of ' x '. The problem, of course, is that while the only difference between the two acts is what is uttered, it does not *seem* plausible to say that what S meant by uttering σ_1 was identical with what S meant by uttering σ_2 .

So we are faced with the following complex dilemma. If by uttering σ_1 S meant that his primary intention . . ., and if S uttered σ_1 with just those intentions purported to be necessary and sufficient for his meaning that p , then not all of the conditions of the account of what it is for S to mean that p are necessary, for the analysans will then exclude at least one case of S meaning that p . Even worse, if in uttering σ_1 S meant that his primary intention . . . and if in uttering σ_2 S did not mean that his primary intention . . . and if the only difference between these two cases is in the means used to make known S 's intentions, then we have at least one case where what is meant is determined by the meaning of the sentence S uttered, and this would seem to call into question the whole enterprise we are engaged in.

One way out would be to deny, *a priori*, that S can utter σ_1 with just those intentions specified in the definition. This line might be taken in either of two ways. (1) It may seem that the

relevant difference between σ_1 and σ_2 is that were S to utter σ_1 his intention would be to produce in \mathcal{A} , by means of recognition of intention, the belief that S uttered σ_1 intending to produce in \mathcal{A} , by means of recognition of intention, the belief that p , whereas if S were to utter σ_2 his intention would merely be to produce in \mathcal{A} , by means of recognition of intention, the belief that p . But we have already seen that this way out is unavailable, since S cannot intend to produce in \mathcal{A} the belief that p by means of recognition of intention without, in effect, intending to produce in \mathcal{A} —by means of recognition of intention—the belief that S uttered x intending to produce in \mathcal{A} the belief that p by means of recognition of intention. (2) It may be argued that σ_1 does not present a counter-example, because one who uttered σ_1 would have *two* primary intentions: one to produce in \mathcal{A} the belief that p , the other to produce in \mathcal{A} the belief that a primary intention in uttering σ_1 was to produce in \mathcal{A} the belief that p . However, this way out may easily be side-stepped by using as our example the sentence ‘My *only* primary intention in uttering this is . . .’ It seems incredible to think that S could not utter a sentence which “explicitly describes” those and only those intentions S had in uttering that sentence.

Alternatively, one may deny, as I propose to do, that in uttering σ_1 S does in fact mean that his primary intention is to produce in \mathcal{A} the belief that p (etc.). Taking this line does not (as noticed in Chapter I) preclude one from saying that by ‘My primary intention . . .’ S meant “My primary intention . . .”

Suppose that there can be found a sentence σ_3 , such that (i) σ_3 means the same as σ_1 , but (ii) it is quite clear that one would want to say that S would *not* mean by uttering σ_3 that his primary intention was to produce in \mathcal{A} the belief that p (etc.). If such a sentence could be found, this would give us extremely good grounds for denying that by uttering ‘My primary intention in uttering this is . . .’ (σ_1) S meant that his primary intention . . .; for it seems highly reasonable to assume that if two sentences have the same meaning, then what would be meant by “subscriptively” uttering the one would also be meant by “subscriptively” uttering the other.

Consider now the explicit performative

σ_3 : I (hereby) tell you that p .

(i) Presumably, as Austin so forcefully urged, in uttering σ_3 one

would not mean *that* one was telling \mathcal{A} that p ; one who utters ‘I hereby tell you that p ’ is not saying *that* he is performing the speech act of telling \mathcal{A} that p . Explicit performatives—sentences such as ‘I (hereby) tell you that p ’, ‘I (hereby) warn you that p ’, ‘I (hereby) promise to ψ ’, ‘I (hereby) request you to ψ ’, and so on—“do not ‘describe’ or ‘report’ or ‘constate’ anything at all, are not ‘true’ or ‘false’.”³ (One might hold that while in uttering σ_3 S is not “constating” that he is telling \mathcal{A} that p , he nevertheless means *that* he is telling \mathcal{A} that p . Such a reply, aside from being false, would be uninteresting; for surely our inclination to say that by uttering σ_1 S meant that his primary intention was . . . is inextricably bound up with our inclination to say that when he uttered σ_1 S said *that* his primary intention was . . .) (ii) Presumably, to tell someone that p is (roughly) to utter x with the primary intention of informing someone—by means of recognition of intention—that p . For instance, if we know that John told Mary that Socrates was a gadfly, then we know that John uttered something intending to inform Mary—by means of recognition of intention—that Socrates was a gadfly. Evidently, then, unless ‘tell’ changes its meaning throughout changes of syntax, ‘I hereby tell you that p ’ (σ_3) has (roughly) the same meaning as ‘My primary intention in uttering this [or, ‘I hereby utter this with the primary intention . . .’] is to produce in you—by means of recognition of intention—the belief (knowledge) that p ’ (σ_1). And so, presumably, what would be meant by uttering σ_1 is identical with what would be meant by uttering σ_3 . And if it is true that in uttering σ_3 S would not mean that he was telling \mathcal{A} that p , i.e., would not mean that he uttered σ_3 with the primary intention of informing \mathcal{A} that p (etc.), then it would seem to follow that in uttering σ_1 S would not mean that his primary intention was to inform \mathcal{A} (or to produce in \mathcal{A} the belief) that p (etc.).

There are two suggestions which may be urged by way of a reply to the preceding argument. One is that in uttering an explicit performative one *does* (*pace* Austin) mean that one is performing a certain speech act, one is “constating” that one is performing a certain speech act. The other is that explicit performative verbs do in fact undergo a “change of meaning” when they occur in the first person singular present indicative active. In the next chapter I offer a detailed account of explicit performa-

³ J. L. Austin, *How to Do Things with Words*, p. 5.

tives, and these suggestions will there be discussed at some length (and dismissed). In the meantime I hope that it will be agreed that a *prima facie* case has been made for thinking that if *S* utters x with just those intentions specified in either definition, then *S* will either mean that p or that *A* is to ψ , but he will not mean anything else, no matter what he uttered.

III.3 *Alleged instances of S meaning that p and S meaning that A is to ψ where the intended response is neither the activated belief that p nor A's doing ψ*

We have now to consider certain alleged threats either to the condition that *S* meant that p only if *S* intended to cause in *A* the activated belief that p or to the condition that *S* meant that *A* was to ψ only if *S* intended to cause *A* to ψ . The allegedly problematic cases are of two types: cases in which the response *S* intends to produce in *A* is neither the activated belief that p nor *A*'s ψ -ing, and cases of utterances in the absence of an audience. The audience-less cases will be considered in section III.4.

(a) Let us begin by considering certain cases which only appear to be counter to the definitions.

(1) It has been objected that one may mean that p or that *A* is to ψ while not caring in the least whether one is believed or obeyed.

Granted that, in uttering x , *S* both meant that p and did not care whether he caused in *A* the belief that p . It does not follow from this that *S* did not utter x with the intention of causing *A* to believe that p .

Customer: "Where is lingerie?"

Clerk: "Lingerie is on the fifth floor."

No doubt the clerk, in saying that lingerie is on the fifth floor, intends to be complying with the customer's request; i.e. the clerk intends to be supplying the customer with the information requested. But it may also be that the clerk is completely indifferent to whether or not the customer accepts this information and that the clerk would not be at all disturbed were the customer to refuse to believe that lingerie is on the fifth floor.

(2) Reporter: "Tell me, Mr. President, what prompted you to lower the draft quota for November?"

Mr. President: "The election." (*Sotto voce*: "I didn't want to let that out!")

Here the speaker "blurts out" something which he intended *not* to divulge. It may seem that while the speaker succeeded in telling *A* that p (and, *a fortiori*, in meaning that p), he did not utter what he uttered intending to inform *A* that p .⁴

I believe that the correct and natural thing to say about this type of case is that *S* uttered x with the *momentary* intention of informing *A* that p , but that this was an intention he intended not to have. That is, *S* forgot his *standing* intention not to inform *A* that p and uttered x with the momentary intention of informing *A* that p . This is not an uncommon phenomenon. One who has given up smoking may, in a moment of forgetfulness, begin to smoke a cigarette. In one sense the smoker intended not to smoke the cigarette, although this was, in another sense, the intention he had in lighting the cigarette. In other words, he had an intention he intended not to have.

(3) There are cases where *S* both means that p (or that *A* is to ψ) by uttering x and utters x believing (or even knowing) that he will not thereby cause *A* to have in mind the belief that p (or to ψ), either because *S* believes that *A* already believes (and has in mind) that p (or has already ψ -ed) or because *S* believes that *A* will refuse to believe that p (or to ψ).

In general, one cannot do an act *X* with the intention of bringing about a certain result if one knows or believes that one will not thereby bring about that result. But while this is generally the case, it is not necessarily the case. A person trapped in a burning building might leap from a seventh floor window with the intention of saving his life. Somewhat analogously, it is not uncommon to try to convince someone that p despite its being virtually certain that one will fail.

(b) However, there are various types of cases of which one would be inclined to say that *S* meant that p (or that *A* was to ψ) even though it is quite clear that *S* had no intention to cause in *A* the activated belief that p (or to cause *A* to ψ).

(1) The Berkeley police, thinking that the campus would be the best place to contain a riot, announce to the Berkeley radicals that under no circumstances are they, the radicals, to hold their rally on the campus. It is the intention of the police that this announcement should cause the rally to be held on campus.

It is clear that the police intend the radicals to *think* that they are being ordered not to meet on campus. One may also feel

⁴ I owe this type of example to Professor P. F. Strawson.

that they *are* ordering this (and so do mean that the rally is not to be held on campus). In this case, as in all cases involving a counter-suggestible audience, *S*'s intention in uttering x is to bring about a certain response by means of deceiving *A* into thinking that *S* uttered x intending to bring about some contrary response.

(2) Teacher: "Tell me, if you can, when the Battle of Hastings was fought."

Student: "1066."

Certainly it is not part of the student's intention to secure that his teacher has in mind the date of the battle. The point of both the teacher's "request" and the student's "reply" is that the student should do what he would do were he actually intending to inform (tell) the teacher of the date of the battle, thereby enabling the teacher to determine whether he does in fact know when the battle was fought. In the "counter-suggestible" example *S* pretended to be uttering x with certain intentions in order to deceive *A* into thinking that he uttered x with those intentions. In the present example *S* openly makes as if to be telling *A* that *p* in order to show *A* that he, *S*, is able to tell someone, and so knows, that *p*.

(3) Police: "O.K., Capone, the jig's up. We know you stole the bubble-gum, so you'd better confess."

Capone: "I confess. I stole the bubble-gum."

Here, unlike the preceding two cases, it is not entirely obvious what *S*'s overt intention is. One may have evidence that *p* sufficient to warrant one's claiming to know that *p*, but if it is important that one be as certain that *p* as possible and if one's evidence is indirect, there may be a point to seeking further evidence. Since Capone is in the best possible position to know what he did, his intention in "confessing" that he stole the bubble-gum may have been to make available the best grounds for entitling the police to claim to know that he stole the bubble-gum. This is especially plausible in the present example, where Capone's utterance may be cited in court as evidence of his guilt.

But there are other "confession" cases which do not so neatly fit this explanation, e.g. George Washington's confession that he chopped down the cherry tree (when it was perfectly obvious that he did). George Washington's "confession" may have been

a gesture to show that he was not dishonest, or that he was not intending to hold anything back.

(4) Mr. Smith: "I was working in the office all evening."
Mrs. Smith: "You're lying."

Apparently, it is not Mrs. Smith's intention to cause her husband to have in mind the fact that he is lying. More likely, her intention is to let him know that she knows that he is lying. Perhaps in such cases one gains a certain emphasis by using a sentence conventionally designed for informing someone that *p* (in this case, that one is lying).

I am inclined to think that none of these cases present a serious problem and that the sense in which it may be said of *S* in these examples that he meant that *p* or meant that *A* was to Ψ is an extended or attenuated sense, one derived from and dependent upon the primary sense captured in the definitions.

(1) "Counter-suggestive utterances". There are several reasons for wanting to say that it is only in an attenuated sense that the police "meant" that the rally was not to be held on campus. First, the very possibility of this type of deception is dependent upon its being *at least* generally the case that one who means that *A* is to ψ intends *A* to ψ . Unless the imperative mood were conventionally correlated with an intention to produce action, it would not be possible to pretend by uttering an imperative that one had such an intention. "Counter-suggestive" utterances, then, are directly parasitic upon primary and standard acts of meaning. Second, there is a tendency to speak of *pretended* or *as if X-ing* as *X-ing*, especially in those cases where what one *does* is exactly the same as what one would do were one actually *X-ing* (in our case, meaning such-and-such). For example, in the primary sense of the expression 'applying artificial respiration', one is applying artificial respiration to another only if one's intention in doing whatever it is that one does is to restore a person to normal breathing, but in a first-aid class the instructor might quite naturally request one student to apply artificial respiration to another student. Third, one will be willing to say that the police *meant* that the rally was not to be held on campus only in so far as one is willing to say that the police *ordered* the rally not to be held on campus. Whatever one's intuitions about

meaning, it seems more difficult to deny that in uttering x S ordered A to ψ only if S uttered x with the intention of getting A to ψ . Finally, it is worth noting that, should it be desirable, the definitions of ' S meant that p ' and ' S meant that A was to ψ ' may easily be altered to accommodate both the primary cases and the "counter-suggestive" cases: we simply require only that S utter x intending A to think S uttered x with those intentions specified in the definitions. But once such degenerate cases have been noticed and accounted for it is probably best to leave the formal definitions unaltered so that they will provide a schema more directly for the standard and primary cases.

(2) "Examination answer". Much of what was said about the preceding type of case applies to the student's "telling" his teacher that the Battle of Hastings was fought in 1066. In the primary sense of 'tell', S told A that p only if S uttered x intending to inform A that p . The student, in uttering '1066', acts as if he were genuinely telling the teacher when the battle was fought, and it is in virtue of this fact that we speak of the student "telling" his teacher that the Battle of Hastings was fought in 1066 and, *a fortiori*, of his meaning that the Battle of Hastings was fought in 1066. (Notice that, in appropriate circumstances, the teacher might have requested the student to "remind" (or "point out" to) him when the battle was fought.)

(3) "Confessions". Certain of these cases may be construed as meeting the conditions for S meaning that p (e.g., where S intends his utterance to be used by others as evidence that p), and others may be provided for by a slight adjustment which allows that S may mean that p when his primary intention is to strengthen an already existing belief that p . (There are analogous imperative cases where S 's primary intention is to strengthen an already existing intention; e.g., S orders A to ψ , which A was intending to do anyway, so that A may have additional reason to ψ .) More problematic are those cases of "confession" where, evidently, S 's intention is neither to create, activate nor strengthen a belief that p . My inclination is to say that here, too, our willingness to say that S "meant" that p derives from the fact that in this type of case (e.g., George Washington's "confession" that he chopped down the cherry tree) S utters x as if he were informing (telling) A that p . Communicating that p is a paradigm of bringing the fact that p "out in the open", and by acting as though he were communicating that he did such-and-such the confessor

purports to show that he is willing to have it out in the open that he did such-and-such.

(4) "Accusations". While it is true that Mrs. Smith meant that she knew (or believed) that her husband was lying, it is, strictly speaking, false that she meant that he was lying. Such cases are not unlike the "examination answer" cases: S shows that he knows that p by uttering a sentence which he would utter were he to inform someone that p .

III.4 Instances of S -meaning in the absence of an audience

These cases divide into two types: (a) those in which S utters x because of the possibility of producing a certain response in some person or type of person, and (b) those cases in which S utters x without having (or without seeming to have) any audience-directed intention at all.

(a) Examples of the first type are:

(i) S , on the off-chance that his mother-in-law will stop by, leaves a note on the door telling her that he will be away for the evening.

(ii) S records in his private diary that his mother-in-law stopped by for the tenth consecutive evening. (Presumably, S feels that his future self may be sufficiently interested in being reminded or told of what transpired in his younger days.)

(iii) Fearing that hippies may discover his land and take to camping on it, S posts a sign saying, "Private property. Keep out".⁵

This range of examples poses three minor problems for the definitions of ' S meant that p ' and ' S meant that A was to ψ '.

(1) The first difficulty arises over the description of S 's primary intention. In example (i), for instance, should we say that S 's primary intention is to inform his mother-in-law that he is away for the evening, or should we only say that his primary intention is to inform his mother-in-law that he is away if she happens by (or, that he left the note intending that, should she happen by, she will be informed that he is away)?

If S utters x with the intention of thereby informing A that p , there are several ways in which he might fail to achieve this end: A might already know that p ; A might refuse to believe that p ; or A might not encounter S 's utterance x . Generally, if S utters

⁵ This example was suggested by an example of Grice's in "Utterer's Meaning and Intentions", p. 172.

x with the intention of informing A that p , S will expect to inform A that p ; but we have already noticed that S may be said to intend to inform A that p even though he may think that A already knows that p and even though he may think that A will probably refuse to believe that p . The case where S utters x thinking that A will probably not encounter x does not seem relevantly different, and so, by parity of reason, the fact that there is a good chance that S 's utterance will go unnoticed does not seem to be sufficient to defeat the claim that he uttered x intending to inform A that p . I think this position is supported by our normal way of describing such cases. For example, we may say that a person's intention in putting a fence around a pit was to prevent someone's falling in, despite the fact that it was known to be highly unlikely that anyone would go near the pit.

The next two difficulties have to do with the mutual knowledge* conditions, and they will require slight revision.

(2) Consider again example (i). Since S will not know if his mother-in-law sees the note, he cannot have intended by putting up the note to realize a state of affairs E which is such that if E obtains, then he and his mother-in-law will mutually know* that it obtains.

But suppose that S 's mother-in-law does encounter the note. She will then know that he knows that the note is on the door; she will then know that he knows that if she knows that the note is on the door, then she will know that he knows that the note is on the door; she will then know that he knows that if she knows that the note is on the door, then she will know that he knows that if she knows that the note is on the door, then she will know that he knows that the note is on the door; and so on, *ad infinitum*.

For this type of case, then, we may say that

S uttered x intending thereby to realize a certain state of affairs E which is (intended by S to be) such that if E obtains, then

- (1a) if A knows that E obtains, then A will know that S knows that E obtains;
- (1b) if A knows that E obtains, then A will know that S knows that (1a);
- (1c) if A knows that E obtains, then A will know that S knows that (1b); and so on.

(That is,

$$\begin{aligned} &K_A E \rightarrow K_A K_S E \ \& \\ &K_A K_S (K_A E \rightarrow K_A K_S E) \ \& \\ &K_A K_S (K_A E \rightarrow K_A K_S (K_A E \rightarrow K_A K_S E)) \ \& \\ &\quad \cdot \\ &\quad \cdot \\ &\quad \cdot \end{aligned} \quad)$$

(3) In example (iii) S 's intention is not to produce a response in a particular person A , but only to produce a certain response in *any person* of a certain type. Consider another example with this feature. S places a "car for sale" notice in a newspaper intending to inform anyone reading the paper of this fact. Obviously, it would be false to say that Smith, who placed the advertisement, and Angleworm, who read it, mutually know* about the advertisement; for it may well be that Smith is completely ignorant of Angleworm's existence. In this case, there is a certain property F such that Smith places the advertisement intending to realize a certain state of affairs E which is such that if E obtains, then

- (1a) if anyone who is F [e.g., anyone who reads the advertisement] knows that E obtains, then that person will know that S knows that E obtains;
- (1b) if anyone who is F knows that E obtains, then that person will know that S knows that (1a); and so on.

A moment's reflection shows that in all instances of S -meaning (so far considered, at least) S 's intention may be so described. For example, S may intend to produce a certain response in anyone who is identical with Angleworm, or anyone who is the one and only man standing immediately before him, etc.

We arrive, then, at the following (quite informal) redefinitions which should accommodate all of the cases considered so far.

S meant that p by 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 to secure that

- (1a) if anyone who has a certain property F knows that E obtains, then that person will know that S knows that E obtains;

- (1b) if anyone who is *F* knows that *E* obtains, then that person will know that *S* knows that (1a); and so on;
- (2a) if anyone who is *F* knows that *E* obtains, then that person will know (or believe)—and know that *S* knows (or believes)—that *E* is conclusive (very good or good) evidence that *S* uttered α with the primary intention
- (1') that there be some ρ such that *S*'s utterance of α causes in anyone who is *F* the activated belief that $p/\rho(t)$;

and intending

- (2') satisfaction of (1') to be achieved, at least in part, by virtue of that person's [i.e., the person(s) satisfying (1')] belief that α is related in a certain way *R* to the belief that *p*;
- (3') to realize *E*;
- (2b) if anyone who is *F* knows that *E* obtains, then that person will know that *S* knows that (2a); and so on.

Likewise for the analysans for the account of what it is for *S* to mean that *A* [or anyone who is *F*] is to ψ .

The definitions on page 63 are seen to be a consequence of the above redefinitions once it is specified that *S* uttered α in the presence of the intended audience, and for all practical purposes the definitions on page 63 will continue to serve as models for the standard and primary case.

(b) We turn now to the alleged cases of *S* meaning that *p* or *S* meaning that *A* is to ψ where there is no audience-directed intention at all.

To begin with it is relevant to recall the distinction made in section I.1 between the two senses in which a person may be said to have meant something. For many of the cases cited as presenting a difficulty for the condition in question are cases where it is true to say that, for example, by "The cat is on the mat" *S* meant "The cat is on the mat", but not true to say that in uttering "The cat is on the mat" *S* meant *that* the cat was on the mat. One might, for some reason or other, while alone and in the privacy of one's study, write out the entire text of *Mr. Apollinax*. No doubt one will have meant something by the words 'He laughed like an irresponsible foetus', but it is unlikely that in writing this one will have meant *that* . . . What we want to find are examples of *S*

meaning that *p* or examples of *S* meaning that *A* is to ψ where *S* has no intention to produce a response in some actual or possible audience (including himself on some later occasion).

Such examples are not, I think, easy to come by. I have neither heard nor (with one possible exception to be mentioned later) been able myself to think of any examples of this type of which it may plausibly be said that in uttering α *S* meant that so-and-so was to ψ . It is not quite as difficult to find candidate examples for the case of *S* meaning that *p*.

(i) A philosopher, alone and in the privacy of his study, is determined to get somewhere on a certain philosophical problem, for example about the relation between meaning and intention. Toward this end he writes down various arguments, suggestions, objections, replies, etc. For example, at one point he writes, "If we should learn that the members of a certain tribe utter the sound 'gavagai' *whenever* a rabbit is present, then the last thing we should say is that in uttering 'gavagai' they mean that a rabbit is present."

Let us suppose that the philosopher has no intention of showing his notes to anyone or of consulting them himself on some later occasion. His intention is to discover the truth.

(ii) A man, during his deliberations about whether to marry Rose, makes a list of her pros and cons. We might suppose that in writing 'She has halitosis' he meant that she had halitosis.

What should we say about these examples?

(1) We might simply deny that we have here instances of *S* meaning that *p* by uttering α . Is it really clear that in (or by) writing 'She has halitosis' *S* meant *that* she had halitosis? Suppose that instead of writing anything down *S* only put together these thoughts in his head. Here it seems to me even less plausible to say that *S* meant that Rose had halitosis when he, as it were, ran the sentence 'She has halitosis' through his mind. But then why should the fact that he did not actually write or utter aloud the sentence constitute a relevant difference?

However, I fear that taking this way out would not win many converts, and for at least this reason I shall henceforth assume that we do want to say of these cases that in uttering α *S* meant that *p*. Another reason for dismissing this hard line is that there is *some* inclination to say that *S* meant that *p*, and this is worth accounting for.

(2) One possible way of accounting for these cases is suggested and illuminated by an analogy with, of all things, jokes. It would appear to be eminently reasonable to think that an utterance-token x is a *joke*—in the primary and important sense of that word—only if someone uttered x intending thereby to amuse someone. But we can easily imagine a dour man, alone in his study, writing ‘My wife is so skinny she has to walk around in the shower to get wet’ with the intention of writing down a joke, and this despite the fact that he has no intention of amusing himself—at that or some later time—or anyone else. Here we are willing to say that S ’s utterance x is a “joke” in virtue of his uttering x with the intention of uttering something which *would* amuse a certain type of audience (if uttered in certain (types of) circumstances).

Similarly, it may be that we are inclined to say of S in examples (i) and (ii) that he meant that p by uttering x in virtue of his uttering x with the intention of uttering something which *would* produce the belief (or activated belief) that p (in the appropriate way, etc.) in a certain type of audience if uttered in circumstances of a certain type. Thus, while in example (i) S does not actually intend to *convince* anyone that p , he may nevertheless criticize the arguments he produces for not being *convincing*.

While an explanation of this sort may account for part of what we have in mind when, of such cases, we say that S meant that p , it cannot be a complete explanation; for there are cases where S utters x intending that his utterance x be such that it would produce the relevant sort of response (in the relevant way) in a certain type of audience, but where it is false that S meant something by uttering x . I do not, for example, think that S meant that p or that S meant that A was to ψ in either of the following two examples.

(iii) A sadistic lieutenant, realizing that he has in his command a naïve and overzealous private, takes delight, while alone, in saying aloud, “Private Goodfellow, run your bayonet through your abdomen, and look sharp about it!”, knowing that were he to utter this in Goodfellow’s presence, Goodfellow would do just that. But it would be a perversion of our intuitions about meaning to say that the lieutenant meant by his utterance that Goodfellow was to run the bayonet through his abdomen.

(iv) A purist, determined never to produce a false sentence, practises on his typewriter by typing the sentence ‘Snow is white’.

Since his intention is to produce a true sentence, his intention, presumably, is to produce something which could be used to tell someone that snow is white. Notwithstanding this, it is false that in typing ‘Snow is white’ he meant that snow was white.

(3) What is the relevant difference between, on the one hand, examples (i) and (ii), where we are inclined to say that S meant that p and, on the other hand, examples (iii) and (iv), where we are not so inclined? I should like to suggest that audience-less examples of types (i) and (ii) are not so very different after all from those cases where there clearly is an intention to produce a response in some audience.

Recall the reasons for thinking that in examples (i) and (ii) S does not intend to produce in anyone the activated belief that p : it is clear that since S already has in mind the belief that p when he utters x , he cannot, in uttering x , intend to produce in himself that response, and since we may suppose that S destroyed what he wrote immediately after his writing session, he cannot have intended the sentence-tokens he produced to cause a response in himself or another on some future occasion.

But consider example (i). S writes down various things intending to arrive at an argument to establish that p . It is not incompatible with the description of this example that S ’s primary intention was to arrive at an argument which could be *reproduced* to convince a certain type of audience that p . Indeed, the academic life being what it is, this may be the most plausible description of S ’s intention. But it would be unduly philistine to think that this is always one’s intention in such cases. One might want to do philosophy even though locked in solitary confinement for life.

But it does seem essential to the “philosophy” example that S ’s intention is to *provide himself* with various arguments, explanations, etc. In this type of case S does not have various arguments, etc., all worked out which he then simply puts on paper (that would be like the “typing” example); rather, he puts various things on paper as part of a process toward arriving at a certain body of *knowledge*. In the audience-present cases of this type S ’s intention is to provide another with certain arguments, explanations, etc.; in the audience-less cases of this type S ’s intention is to provide himself with certain arguments, explanations, etc.

Example (ii) is the audience-less analogue of the "fact review" example described in section II.3. There *S* uttered 'Rose earns 30 pounds a week . . .' with the intention of securing that *A* has all of the relevant facts about Rose in mind while reaching a decision as to whether he should marry her. In example (ii) *S* is doing this job for himself.

I submit, then, that in each case where it is both the case that we are inclined to say that *S* meant that *p* by uttering *x* and that *S* apparently had no audience-directed intention, *S*'s utterance *x* will be part of some activity directed towards securing some cognitive response in himself, and that it is in virtue of this significant resemblance to the standard case that we class these cases as instances of *S*-meaning.

I mentioned above that it is extremely difficult to find examples of the relevant sort where we would be inclined to say that *S* meant that so-and-so was to ψ . The following example may be an exception. Suppose I am planning a bank robbery. I want to secure that I have a complete plan in mind when I undertake this endeavour, and towards this end I write down, 'First, knock out alarm system; second, shoot guard . . .'. Here it does not seem too unnatural to say that I meant that I was first to knock out the alarm system, etc., and this is so because my intention in writing what I wrote was to secure that I would have in mind a network of intentions.

The proffered explanation of audience-less cases also accounts for why it is difficult to find examples where we would be willing to say that *S* meant that *A* was to ψ and why we are not inclined to say that *S* meant something by uttering *x* in examples (iii) and (iv).

III.5 *S* meant something by uttering *x*

Finally, there remains to be defined '*S* meant something by (or in) uttering *x*'. Two complementary arguments will be given to show that *S* meant something by uttering *x* if and only if, for some *p*, *A* [or *F*], and ψ , by uttering *x* *S* meant that *p* or *S* meant that *A* [or anyone who was *F*] was to ψ . (The qualification in brackets should be understood throughout.)

The first argument is designed to show that if *S* utters *x* with the intention of thereby producing a certain response *r* in an audience *A*, then *S* will mean something by uttering *x* *only if*, for some *p* and some ψ , the response *r* is *A*'s (actively) thinking

that *p* or is *A*'s doing ψ ; i.e., if *S* utters *x* with the primary intention of producing in *A* a certain response *r* and if the value of '*r*' is anything other than a belief or action, then *S* will not have meant anything by uttering *x* (unless, of course, *S* also uttered *x* with a primary intention to produce in *A* the activated belief that *p* or a primary intention to get *A* to ψ). If this can be shown, then, given that to mean something *S* must intend to produce some response in an audience and given the correctness of the definitions of '*S* meant that *p*' and '*S* meant that *A* was to ψ ', it will follow that *S* meant something by uttering *x* if and only if, for some *p*, *A*, and ψ , by uttering *x* *S* meant that *p* or *S* meant that *A* was to ψ .

The second argument proceeds primarily by way of showing that if *S* utters a well-formed sentence of English with that sentence's full conventional force, then, for some *p*, *A*, and ψ , *S* will mean thereby either that *p* or that *A* is to ψ .

First argument

Consider the following cases.

(i) *S* snubs *A* in the street, intending thereby to cause *A* to be distressed. (This example is the only explicit reason given by Grice for allowing affective attitudes into his account of *S*-meaning: "if I cut someone in the street I do feel inclined to assimilate this to cases of meaning_{nn}, and this inclination seems to me dependent on the fact that I could not reasonably expect him to be distressed (indignant, humiliated) unless he recognized my intention to affect him in this way."⁶)

(ii) *S* does the washing-up to please his wife. Since his doing the washing-up will please her only if it is done to please her, he intends her to recognize his intention to please her.

(iii) *S* utters *x* with the intention of causing—by means of recognition of intention—*A* to be afraid of a certain dog, Fido.⁷ More specifically, *S* intends *A* to recognize *S*'s intention to get *A* to fear Fido and to have this as his reason for being afraid of Fido.

If feeling distressed, pleased, or afraid were relevant responses for an account of *S*-meaning, then it would be the case both that (1) *S* meant something in one or more of the examples (i)–(iii) and (2) *what* *S* meant [for convenience, Φ] is such that *S*'s intention to cause *A* distress (pleasure, or fear)—in a certain way,

⁶ "Meaning", p. 384.

⁷ By using '*x*' rather than some sentence I am being charitable.

etc.—is either a sufficient condition for S meaning Φ or a necessary condition for S meaning Φ (or a necessary part of a sufficient but not necessary condition for S meaning Φ , etc., although this much caution will not be necessary). (Clearly, it is not enough that S meant something in each (or any) of the above examples. First, if S meant something, then there is (as it were) something that S meant. Second, it may be that while S meant something by uttering x he meant something as a result, say, of his intention to produce in A the belief that p and not in virtue of his intention to cause A to be afraid. For example, a second might shout “Get angry!” to his lugubrious prize-fighter, and here the second would have meant that the prize-fighter was to make himself angry; but he would have meant this in virtue of his intention to cause A to *make himself* angry and not in virtue of his intention to produce the response of anger.)

What, if anything, does S mean in example (i)? A plausible answer is that S meant that he was through with A (this is in fact the answer given to me by Grice when I asked him this question). But if the definition of ‘ S meant that p ’ is correct, then the only response relevant for S meaning that he is through with A is A thinking that S is through with him. Example (i) gives us no reason to doubt this result, for S could not have intended to cause A distress unless he intended to cause A to think that S was through with him. (It may be that A would not think that S intended to produce in A the belief that he, S , was through with A unless A thought that S intended to cause A distress, but that is irrelevant.)

Similar conclusions are reached with regard to examples (ii) and (iii). If S meant anything by washing-up it was that he wanted to please his wife, or that he was doing the washing-up in order to please her. Likewise, in example (iii) S either meant that Fido was dangerous or that there was reason to fear Fido or that A should be on guard against Fido or, conceivably, that A was to make himself afraid of Fido. But no matter which of these things S meant, it is both necessary and sufficient (given, of course, satisfaction of all relevant conditions) for S meaning what he meant that the intended response be A ’s (actively) thinking that p or else A ’s doing ψ .

Should the preceding be correct, we may conclude that neither distress nor pleasure nor fear is a relevant response; and there is reason to think that this conclusion may be extended to

encompass all affective attitudes and emotions. First, I submit that if the reader considers any example in which S utters x intending to produce an affective response in A , then, should S mean anything by uttering x , the reader will find either that, for some p , S means that p or that, for some ψ , S means that A is to ψ (at least this has been my experience). Secondly, it seems implausible that what should apply to distress, pleasure, and fear with regard to meaning should not apply as well to any other affective attitude or emotion. And since other types of propositional attitudes—e.g., doubting, suspecting—are definable in terms of belief or intention, we seem, as intended, to be left with only belief and action (intention) as the only relevant types of response.

If it is true that one can mean something only in virtue of an intention to produce (activated) belief or an intention to produce action, then there should be some important difference between believing and intending and all other types of responses for which one may have reasons (viz., affective attitudes and some emotions). I believe that the relevant difference is one noticed by Grice in “Meaning”. Grice pointed out that “to have a reason for believing so-and-so . . . is . . . like ‘having a motive for’ accepting so-and-so”, and that decisions “that” seem to involve decisions “to”, but that “one cannot in any straightforward sense ‘decide’ to be offended” (p. 386). Thus the relevant difference between beliefs and actions, on the one hand, and affective attitudes on the other may be that the elements of deliberation and decision are absent in the case of affective attitudes. When all the facts are known, one may deliberate about whether or not one should believe that p or about whether or not one should ψ . But when all of the relevant facts are known one either hates Angeworm or one does not. It would be most odd for one to ponder, “Well, Angeworm slandered my wife, is generally odious and malicious. Should I hate him?” (There is a sense in which one might ask “Should I be angry with him?” But such a question most naturally arises, e.g., when one isn’t sure how one is to take what someone said or did, and this is a case where all of the relevant facts are not known.)

Second argument

Let ‘ σ ’ stand for any well-formed English sentence. I want to say that if S utters σ with its full conventional force, then, for some p , A , and ψ , what S will mean by (or in) uttering σ is

either that p or else that A is to ψ . Before arguing for this a few preliminary comments will be useful.

(1) I shall use the word 'sentence' in such a way that 'The cat is on the mat.' is a sentence, but 'the cat is on the mat' is not a sentence. (In speech this distinction is marked by pauses, intonation, linguistic context, etc.)

(2) "Full conventional force." For example, S utters 'The cat is on the mat.' with its "full conventional force" (roughly) if and only if S utters 'The cat is on the mat.' and means thereby that the cat is on the mat. If, say, S utters 'The cat is on the mat.' and means "The cat is on the mat" by 'The cat is on the mat.' (as has been true of me in using this sentence as an example) but does not mean by uttering 'The cat is on the mat.' that the cat is on the mat, then S has not uttered 'The cat is on the mat.' with its full conventional force. Loosely speaking, to utter σ with its full conventional force is to utter σ and to mean thereby what one says. So, again loosely speaking, I want to say that if S utters a well-formed English sentence and means what he says, then, for some p and some ψ , he will either mean that p or else he will mean that so-and-so is to ψ . (For convenience, I ignore problems of ambiguity.)

(3) I do not maintain that if asked what S meant by uttering σ one must reply by *saying* that he meant that p (or that A was to ψ). One might reply in all sorts of ways: "He said that p "; "He asked whether p "; "He meant the cat was on the mat"; etc. However, I do want to say that in so far as these replies are genuine replies, then they will entail that S meant that p (or that S meant that A was to ψ).

The argument is quite simple.

(i) If σ is a well-formed English sentence, then (roughly but not too roughly speaking) σ will be in the indicative, subjunctive, optative, imperative, or interrogative mood. (This is a more generous use of 'mood' than most grammarians would allow, but that is irrelevant.)

(ii) If σ is in the indicative mood and if S utters σ with its full conventional force, then, for some p , S will mean that p by (or in) uttering σ (and, in the absence of further intentions, S will not mean anything else by uttering σ ; this rider should be understood throughout). For example, if S utters 'It is raining.' with its full conventional force, then S will mean that it is raining. On the whole this is obvious enough, but there may seem to be

a difficulty with explicit performatives. However, in Chapter IV I shall argue that if, say, S were, *per impossibile*, to utter 'I request you to pass the salt.' with its full conventional force, then he would mean thereby that he is requesting A to pass the salt. (As suggested in section III.2, meaning being what it is, S cannot mean that he is requesting A to pass the salt.)

(iii) If σ is in the subjunctive mood and if S utters σ with its full conventional force, then, for some p , S will mean that p by (or in) uttering σ . For example, if S utters 'If the cat had been on the mat, then hairs would not be on the floor.' and means what he says, then S will mean that if the cat had been on the mat, then hairs would not be on the floor.

(iv) If σ is in the optative mood and if S utters σ with its full conventional force, then, for some p , S will mean that p by (or in) uttering σ . Thus, if S utters 'Would that my own true love were in my arms!' and means what he says, then he means that he wishes that his own true love were in his arms. (Such a sentence, if uttered at all, would perhaps be uttered as a "pure expression of feeling"; but in this case σ would not have been uttered with its full conventional force and S would not have meant anything by uttering σ .)

(v) If σ is in the imperative mood and if S utters σ with its full conventional force, then, for some A and some ψ , by (or in) uttering σ S will mean that A is to ψ . For example, in uttering 'Shut the door!' S meant that John was to shut the door.

(vi) If σ is in the interrogative mood and if S utters σ with its full conventional force, then, for some A and some ψ , by (or in) uttering σ S will mean that A is to ψ . Questions, I want to say, are requests for information. So if I ask you what time it is I am requesting you to tell me the time. If I ask you whether or not snow is white, then I am requesting you to inform me of whether or not snow is white. If we were to analyse questions in terms of force or mood indicating device plus propositional content or sentence radical, then we would (essentially) represent 'What time is it?' as

! [your informing me of the time];

and we would (essentially) represent so-called "yes-no" questions, such as 'Is it raining?', as

! [your informing me of whether or not it is raining].

I believe that apart from any particular theory of *S*-meaning this account of interrogatives is plausible; but—taking into account the principle that if the only difference between two utterances is that in one case *S* utters x and in the other y , then what *S* meant by uttering x was identical with what *S* meant by uttering y —we have a special reason: if I utter ‘What time is it?’ with its full conventional force, then I will utter it with (roughly) the intention of thereby getting *A*—by means of recognition of intention—to inform me of the time, and if I utter ‘Inform me of the time, please.’ with its full conventional force, then I will utter it with (roughly) the intention of getting *A*—by means of recognition of intention—to inform me of the time; and since the only difference between the “question” and the “request” is in the type of sentence uttered (the means used to make known my intentions), there will be no difference in what I meant.

So we have the conclusion that if *S* utters σ and means what he says, then, for some p , *A*, and ψ , by (or in) uttering σ *S* will either mean that p or else he will mean that *A* is to ψ . This is not yet enough to give us the conclusion that *S* meant something by (or in) uttering x if and only if, for some p , *A*, and ψ , by (or in) uttering x *S* meant that p or *S* meant that *A* was to ψ . We will get that desired result with the addition of the following premiss, which, strictly, is more than is needed.

Let ‘ Φ ’ be a “dummy” for any expression which can complete the sentence form ‘*S* meant that . . . by (or in) uttering x ’. If by uttering x *S* meant that Φ , then there is some (English) sentence σ such that if *S* utters σ with its full conventional force, then *S* will mean that Φ by uttering σ . (This premiss expresses the belief that whatever can be meant can be said. I do not know how to prove this assumption, but given the general thesis about meaning and language being proffered, it is not difficult to think of reasons why this premiss should be true.) This, together with the premiss that if *S* utters σ with its full conventional force, then, for some p , *A*, and ψ , by uttering σ *S* will mean that p or else *S* will mean that *A* is to ψ , yields the conclusion that *S* meant something by uttering x if and only if, for some p , *A*, and ψ , by uttering x *S* meant that p or *S* meant that *A* was to ψ .

There is one consequence of this definition of ‘*S* meant something by uttering x ’ worth mentioning. It is an essential part of

the Austin–Wittgenstein legacy that the class of utterances does not neatly and simply divide into two jointly exhaustive and mutually exclusive subclasses: “statements” and “imperatives”. If what I have said in this section is correct, there is very good reason to think that this legacy is false. More will be said about this in the next chapter.