

How do Gestures Succeed?

Jonathan Bennett

[From in E. LePore and R. Van Gulick (eds), *Searle and his Critics* (Cambridge, MA: Blackwell, 1991), pp. 3–15.]

1. Meaning and intention: why it matters

If we human beings are to command a whole, clear view of ourselves, we need to understand those aspects of our nature that we share with many other kinds of animal and the aspects that are, so far as we know, unique to ourselves; and we need to be able to hold the two sets of aspects together in a single, coherently integrated picture. Perhaps the biggest single obstacle to our doing this is language. It is so conspicuously unique to our species that it tends to impede our view of anything else that is special about us, and yet it is so pervasive and familiar that in a certain way it tends to drop out of sight. If we are to be able to see it as a separable but pervasive part of the human condition, we need to get it in focus—to stand back from it, and from ourselves, just far enough to see how our language relates to the rest of us. Paul Grice's paper 'Meaning', with help from two subsequent papers in which he developed the central idea there presented, helps us to do this and thus to see ourselves whole and clear.¹

In these papers Grice explains the concept of meaning in psychological terms: an analysis of 'By making that gesture,

Joe means that he is hungry' in terms of what Joe intends to achieve by making the gesture. Once we are straight about how *meaning* fits into the human condition, there is no insuperable obstacle to bringing *language* into the picture; for language is basically and essentially and almost exclusively a systematic vehicle of meaning, and the concept of *system* is one that we understand. All that remains, if Grice is right, is to ground intention in our deeper natures. All sorts of people, functionalists and others, are working on the project of explaining intentional psychology in terms of biology, and there are good grounds for optimism. In short: we have a link from biology to psychology, another from that to meaning, and a third from that to language. When these are all worked out properly, that will be a great achievement—one of the greatest achievements that could be expected of philosophers. Grice's contribution is the forging of the middle link, which elucidates the concept of meaning in terms of more general psychological terms.

2. Grice's theory

It is plausible to think that meaning is essentially tied to communication, so that what someone means by a noise of

¹ H. P. Grice, 'Meaning', *Philosophical Review* 66 (1957), pp. 377–88; 'Utterer's Meaning, Sentence-Meaning, and Word-Meaning', *Foundations of Language* 4 (1968), pp. 225–42; 'Utterer's Meaning and Intentions', *Philosophical Review* 78 (1969), pp. 147–77.

gesture may be explained or analyzed or defined in terms of what he intends to communicate to his audience. The general idea breaks down into two main cases. The person intends **(1)** to get his audience to believe that P, or **(2)** to get his audience to do A, where A is some kind of action. In **(1)** a *statement* is made, in **(2)** an *injunction* is uttered—using ‘injunction’ to cover commands, requests, bits of advice, and any other utterances whose immediate purpose is to influence the audience’s behavior.

I give pride of place to the split between statements and injunctions because it is deep and central. Kinds of speech act that belong to neither of these categories, though they contribute greatly to the richness of our scene, are not necessary to language as such, are explicable in terms of statements and injunctions, and should be introduced at a later stage. For now I shall confine myself to statements, but everything I shall say can easily be adapted for injunctions as well.

The suggestion is that ‘By doing x, S means that P’ means about the same as ‘By doing x, S intends to get someone to believe that P’. It has long been known that our ordinary notion of *meaning that P* is not captured by this analysis, because we often do things intending to get others to believe that P without meaning that P by what we do. For example, I might show you a photograph of McStiggins streaking at a football game, intending thereby to get you to think that he has behaved in that way; but in showing you the photo I don’t *mean that* McStiggins streaked at a football game. Showing you that photo isn’t remotely like making a statement.

On the other hand, if I show you a *drawing* of McStiggins streaking at a football game, intending to make you think that he has behaved like that, that is a more plausible candidate for meaning. What is the difference? It is that the photograph’s status as evidence that P has nothing to do

with why I showed it to you, whereas the drawing can serve for you as evidence about McStiggins’s behavior only if you think I showed it to you for that purpose. If in the second case I sanely show you the drawing as a way of getting you to think that (P) McStiggins streaked at a football game, I must be intending that you take it as evidence that P; but that could only be *intention-dependent* evidence; and that, according to Grice, is what marks off meaning from other kinds of belief-production.

That is how I express Grice’s theory. He puts it like this: In doing x, S means that P if and only if:

In doing x (1) S intends to get an audience to believe that P, and (2) S expects this to come about through the audience’s realizing that (1) is the case.

A similar story holds for injunction. Compare these two cases in which I do something intending to get you to close the window. **(a)** I act the part of someone suffering from a draft on the back of his neck (or I show that I *am* suffering from a draft); I want you to take this as evidence that the draft is bothering me, and to close the window. **(b)** I go through a mock performance of someone suffering from a draft, and then make a big to-do of looking at the window. Nobody could possibly take my shivers and neck-rubs as natural, uncontrived expressions of suffering from a draft, and I don’t expect you to do so. Rather, I want you to take them as evidence that I want you to close the window, and to be led by that to close the window. How can I expect you to take such a transparent charade as evidence that I want you to close the window? By taking it that I am offering the charade for that purpose. So the core of this injunction, as of the statement, is the attempt to give intention-dependent evidence for something.

3. Meaning-nominalism and other complaints

A great feature of this account of the concept of meaning is that it helps us to see that meaning doesn't necessarily involve rules or conventions. Uttering *x* and thereby meaning that *P* can be a one-shot deal, though of course there can be conventions for meaning (and a language is a system of such conventions). Analogously, we could have a convention for driving on the right, but the notion of driving on the right can be fully understood without bringing in the concept of convention. One of the things I did in my book *Linguistic Behaviour* was to show how to take the concept of convention (as analyzed by David Lewis) and the concept of meaning (as analyzed by Grice), to hold them apart and then bring them together in a fruitful combination.

Some of the earliest attacks on Grice's work on meaning focused on its 'meaning-nominalist' aspect, i.e. on its treating meaning as something that can occur in an isolated case that owes nothing to rules or regularities or conventions. None of these attacks was convincing, because they were all backing a loser. An early paper of Searle's, a modified version of which is in his book *Speech Acts*, purported to show that Grice's analysis was wrong because, as Searle put it, 'Meaning is more than a matter of intention, it is also a matter of convention'.¹ But Searle's attempt to show this relies on a discussion that becomes vague at crucial points, and on a supposed counterexample to Grice which I have argued (to my own satisfaction) doesn't work.²

Those who hold that meaning requires convention must conclude that Grice's analysis is too weak to be sufficient

for meaning. Strawson has given another reason for the same conclusion; but he did not offer it as fatal, because he rightly thought that the trouble could be repaired without detriment to the main thrust of the analysis.³

That's enough about the charge that Grice's analysis is too weak for meaning. My main topic comes from the correct accusation that the analysis is too strong.

My purpose in saying to you that *P* need not be to convince you that *P*; it may instead be to show you that *I* know that *P*, or to *remind* you that *P*. We can easily weaken the account of what the 'speaker' intends enough to take in these two possibilities: replace 'intend to get the audience to believe that *P*' by 'intend to get the audience to believe that the speaker believes that *P*' or 'intends to get the audience to bring actively to mind the belief that *P*' or the like.

Something along these lines might be made to yield the view—which some have found plausible—that in statement-making the primary aim of the speaker is to produce *understanding* in the hearer. Of course that would be no help in an analysis of meaning if 'y understands x' has to be explained as 'y knows what x means'; but it may be possible to bring understanding into the picture while avoiding that circularity.

4. Absent audience

Amendments of the kind I have described so far leave the spirit of the analysis untouched. But other troubles are more recalcitrant. According to a recent paper of Searle's, one may speak to someone in a statement-making kind of

¹ John R. Searle, 'What is a Speech Act?', in M. Black (ed.), *Philosophy in America* (Cornell University Press: Ithaca, NY, 1965), pp. 221–39, p. 230.

² John R. Searle, 'What is a Speech Act?' and *Speech Acts* (Cambridge University Press, 1969), pp. 47–50; Jonathan Bennett, 'The Meaning-Nominalist Strategy', *Foundations of Language* 10 (1973), pp. 141–68, pp. 164f.

³ P. F. Strawson, 'Intention and Convention in Speech Acts', reprinted in his *Logico-Linguistic Papers* (Methuen: London, 1971), pp. 149–69. The best way to meet Strawson's point, I believe, is that presented in my *Linguistic Behaviour* at pp. 126f.

way without caring what effect one has on the person. He writes:

'Why can there not be cases where one says something, means it, and does not intend to produce understanding in the hearer? In a case, for example, where I know that my hearer is not paying attention to me I might feel it my duty to make a statement even though I know he will not understand me. In such cases the speech act is indeed defective, because the speaker fails to secure illocutionary up-take; but even in such cases it seems clear that the speaker means something by what he says even though he knows his speech act is defective.'¹

That seems right, and there are other kinds of example as well. Someone who speaks to himself does not even have an audience, yet his use of language is presumably meaningful. It isn't a mere accident that he mutters 'I wish I were dead!' rather than 'Oh, let it last for ever!'

The morning after I discussed this subject in public for the first time, the morning newspaper carried a comic strip in which the egregious Andy Capp looks carefully to left and to right, sees that the bar is empty except for the barman who is dozing, and shouts 'Drinks all round!' He then turns to the reader and confides: 'I've always wanted to say that.' That final frame forces the point home: when Andy Capp says 'Drinks all round!', it's not just that his words mean something—he means something by them. So Grice's analysis is too strong.

In *Linguistic Behaviour*, which amends and develops and

builds on Grice's theory of meaning, I proposed that we to stop trying to formulate conditions that are not only sufficient for meaning but also necessary for it:

'If a community's verbal behaviour was indistinguishable from that of the characters in Chekhov's *The Three Sisters*, say, that might suffice to qualify it as linguistic; but such a specialized example does not illuminate the concept of meaning or of language. What we need, at a minimum, is a statement of conditions which are strong enough to be sufficient for meaning and yet *weak enough to be instructive*. . . [That] restores the problem to the form: how can we get something strong enough but not too strong?'²

My point was that the search for a biconditional analysis is a search for conditions that are strong enough to be sufficient for the analysandum yet *weak enough to be necessary* for it; and I was proposing that we give up looking for this much weakness, and settle for something that is present throughout a subclass of instances of meaning that is large and central enough to serve as a good basis upon which to understand all the cases of meaning that lie outside it. I think of it as a secure base-camp from which one can make brief forays into the surrounding territory.

5. Explaining the absent audience cases

Virtually all the cases of meaning that stubbornly resist Gricean analysis—such as Andy Capp's offer of drinks—have in common that they involve language, i.e. systematic conventional meaning. I contend that if we start with meaning

¹ John R. Searle, 'Meaning, Communication, and Representation', in R. Grandy and R. Warner (eds), *Philosophical Grounds of Rationality: Intentions, Categories, Ends* (Oxford University Press, 1986), pp. 209–26, p. 211.

² Bennett, *Linguistic Behaviour*, pp. 22f. The same device deals also with some difficulties that could be handled within the Gricean framework. They are so handled by Stephen R. Schiffer, *Meaning* (Oxford University Press, 1972), but the result of thus soldiering on with biconditionals is that one ends up with an analysis that is not useful because it is too heavy to lift.

as an isolated phenomenon, we shall encounter hardly any such cases. So we can use Grice's analysis for the notion of meaning as an isolated phenomenon, then bring in *convention*, move from that to *language*, and then be in a position to open the floodgates to all the fancy meaningful things we can do because we have language. (I cannot prove this; I'm offering it as a plausible conjecture.)

What, then, should we say about Andy Capp's offer? It is obviously true and unproblematic for Gricean theory to say that he utters a sentence whose conventional meaning is an offer of drinks to all the members of a present audience. But what makes it the case that *Andy* means something by the sentence (or by uttering it)? I suggest that it is this fact:

A full explanation of why he utters that sentence rather than some other will involve some fact about what he takes the conventional meaning of the sentence to be.

And what makes it true that Andy means the sentence as an offer of drinks is this:

A full explanation of why he utters that sentence rather than some other will involve the fact that he takes that sentence to mean conventionally an offer of drinks.

Something along those lines will probably serve for all the other cases that are not enclosed within the base camp, not captured by the initial Gricean story.¹ Some of them will have flavors of their own, coming from special links between the speaker's mind and the conventional meaning of what he utters; it is not to be expected that all this could be expressed in a unitary account. But the above treatment is the part of the story that is common to all the cases, and will be most of the story in each individual case.

¹ It also fits all the Gricean cases, where real audiences are genuinely offered drinks. The two classes come apart in their further details: they need different accounts of how the conventional meaning of the sentence figures in the explanation of why the speaker utters what he utters.

6. A pocket of resistance

John Bricke has suggested to me that some cases of 'no intended-audience' meaning don't involve meaning conventions. I sit in my study replaying in my mind a recent conversation with a colleague; it suddenly dawns on me that I have said something stupid or harmful, and that the damage cannot be undone; and I make a slashing movement across my left wrist with my right forefinger, meaning by this something like 'I'm so angry with myself that I could kill myself'. The case is not offered as one in which that is literally *true*, just as one in which it is *meant*; similarly, one might utter those words meaning something by them yet not speaking truly.

For this example to embarrass my Gricean program, it must involve 'non-natural meaning', as Grice calls it. That is, my gesture must be something by which *I* mean that I am angry etc.; it mustn't be the case merely that my gesture 'means that I am angry etc. in the way that my sniffing 'means' that I have a cold. But I find it intuitively plausible to suppose that the gesture is meaningful in the former way, which makes this a problem for my program.

I allowed for this by saying that 'virtually all' rather than 'all' *simpliciter*, the no-intended-audience cases involve language. What can I say about the ones that don't? If there are any, i.e. if cases like Bricke's really do involve non-natural meaning, then the sorties from the Gricean base-camp will not all use the tactic of explaining what the speaker means in terms of how his mind relates to what his utterance conventionally means. I could live with that loss of unitariness in the procedure, but I don't know what other procedure(s) to use for the recalcitrant cases. I find

it reasonable to suppose that the person who makes the finger-across-wrist gesture may mean that P by it, for a certain P; but I don't yet know how to justify or explain that intuition on the basis of any general truths about meaning.

7. Representation to the rescue?

Despite this unsolved problem, I think that my base-camp strategy is reasonable and principled, not a mere ad hoc retreat in the face of the no-intended-audience problem. It may be the best that can be done. Given the great variety and complexity of things we can do with language, it would be surprising if we could capture the phenomenon of meaning that lies at their core in a single unitary analysis.

Still, it would be nice if we could: the old idea of necessary and sufficient conditions is still attractive, if only we could implement it for the concept of meaning.

Searle suggests that we can. This is in the recent paper of his, in the Grandy and Warner volume mentioned above, in which he denies that someone who means must intend to affect an audience. Searle argues there that Grice has given us a theory not of meaning but of communication; that the characteristically Gricean flourish—the piece about intending to affect someone through his realizing what you are up to—is true of communication; and that meaning can be analyzed in terms that don't involve any thought of an audience.

According to Searle, for someone to mean that P by uttering x is for him to utter x *intending that it represent the state of affairs that P* ('Meaning, Communication, and Representation', p. 216). This looks promising: it leaves the audience out of the picture, while making it easy for it to be brought into it: when I do something intending it to represent the state of affairs that P, I may intend further that my audience recognize this first intention and be led by it

to believe that P, or to bring it about that P. Searle shows that if you handle meaning in his way, and then take the most natural further step to communication, you will end up with an account of the latter that really is Gricean; so Grice is allowed his success, not at the ground floor where he thought it was taking place, but one level up, as a natural consequence of Searle's account of the ground floor.

This bypasses the worst obstacles to a Gricean biconditional about meaning by offering only a Gricean biconditional about communication, with this resting on a non-Gricean, Searlean biconditional about meaning. This is neat and pleasing and plausible. It would be hard to resist *if* one were satisfied that the notion of *representation* could be adequately explained without help from any such notions as those of meaning and communication. For example, it might be that 'x represents the state of affairs that P' means something like 'Somebody produces x as part of an attempt to get someone else to have the thought that P' or, worse still, 'somebody produces x thereby meaning that P'. If we can't explain *representation* except in such terms as those, Searle's account is vitiated by circularity.

8. What is representation?

The first step in Searle's account of representation says that for x to represent that P, someone must produce or use x with the intention that it represent that P. This implies that Searle's basic analysandum is not *x represents that P* but *S produces or uses x intending it to represent that P*; and the circularity problem then arises in connection with the latter.

Searle knows it, and expresses himself cautiously about his chances of success:

What exactly do I intend... when I intend that [x] represent [that P]? The answer to that question, if it is to have any real explanatory power, should be given

in terms that employ no such semantic notions as reference, truth, meaning, propositions, etc. Perhaps, in the end, it will prove impossible to give an answer that does not employ such notions. However, at least a start on such an answer can be made by showing how representing relates to intentional behaviour generally. ('Meaning, Communication, and Representation', pp. 214f)

He confidently offers a conditional of the form 'If (representation) then (such and such)', apparently having doubts only about whether its converse is true. He seems to think that it may be, for he more than once says that representation consists 'at least in part' in his necessary condition, hinting that the latter may be sufficient too. Shortly after that, Searle speaks of 'my account of representation', suggesting that he thinks the account is complete, or probably or nearly so.

If none of these is true, then why should we think his account can be completed without circularity or regress? And if we don't think that, why should we regard this as a significant rival to the Gricean base-camp strategy?

The question of whether Searle's account is complete or at least completable without circularity would be crucial if the account were correct as far as it goes. But it is not. What Searle offers is not even necessary for representation. For the remainder of this discussion, the topic will not be partialness but untruth. Here is Searle's supposed necessary condition for representation:

'Whenever S produces x with the intention that it represent a state of affairs A then it must be the case that S produces x with the intention that a criterion

of success of his action should be that A obtains, independently of the utterance.'¹

In discussing this, I shall assume that 'He intends that a criterion of success of his action should be that P' means something of the form

'He intends it to be the case that: he performs an action for which a criterion of success is P',

rather than of the form

'He intends it to be the case, with respect to the action that he performs, that a criterion of success for it be P.

What is being said is not that he intends to confer upon a certain action a certain relational property; rather, something is said about what kind of action he intends to perform. If this is wrong, I am lost right at the outset.

Throughout the next three sections, I shall further assume that 'A criterion of success for what x intends to do is P' entails 'If not-P then x fails in what he intends to do'. Searle would reject that entailment, I think, but it matters to see that he *must* reject it; in section 12 I shall consider what notion of 'condition of success' he might use that would not have that consequence.

9. Representation and truth

Now, to understand Searle's conditional, we must understand its phrase 'success of his action'. I think Searle means by it 'his success in using x to represent that P', so that the conditional entails that *S hasn't succeeded in using x to represent that P unless P*. That reading seems to be implied by this:

¹ Searle, 'Meaning, Communication, and Representation', p. 215. How does this 'show how representing relates to intentional behaviour generally'? Only by implying that when we intend to represent, as when we intend to do anything else, we may succeed or fail. This isn't much of 'a start' towards analyzing the concept of representation in a way that would fit it to support a theory of meaning. Furthermore, there is reason to think that Searle would not claim even that thin connection; for his notion of 'conditions of success' seems in the event to have nothing to do with succeeding or failing.

'My intention to represent can at least in part be analysed as the intention that certain conditions of success of the utterance be satisfied. How exactly can we specify these? The most obvious way would be to ask what counts as a failure. When would we say the speaker had made a "mistake"? In the case of the broken crankshaft an obvious mistake would be made if S's crankshaft were not broken. If S's crankshaft is not broken then he has made a mistake in representing it as broken.'¹

Something similar happens here:

'The difference between saying it is raining and meaning it, and saying it without meaning it can be got at by examining the question, what counts as a mistake? What counts as a relevant objection? If I say "It's raining" and mean it, then if I look out of the window and see that the sun is shining and the sky is blue, I am committed to recognizing these states of affairs as relevant to my utterance. . .'. ('Meaning, Communication, and Representation', p. 217)

Well, 'relevant' is safely weak; but my concern is with Searle's use of the stronger words 'mistake' and 'objection'. Judging by these two passages, and by his 'criterion of success' conditional, Searle is evidently taking representation to be *true* representation.

This is too narrow a base upon which to erect a theory of meaning. In Searle's crankshaft example, the protagonist does think his crankshaft is broken, and has intentions which cannot be realized if the crankshaft is sound. But there are plenty of cases of communication (and thus of meaning (and thus of representation)) where S intends to produce a false belief in the mind of the audience, by using x

to represent that P when in fact not-P. In such a case, S has succeeded in representing that P, in every ordinary sense of 'represent', and he may well mean that P as well; yet this lies outside the purview of Searle's account.

Incidentally, Searle would not be helped by a weaker reading of his conditional according to which its consequent merely says that S intends to be doing *something* such that: if not-P then he fails in it. If S is lying, he may represent (and mean) that P when really not-P, and may succeed gloriously all down the line, failing in nothing that he intends to do.

10. A base-camp strategy for Searle?

Searle might reply that non-liars are central and basic to his analysis of meaning, and that he plans to build his account on them and then move out to capture the liars. This resembles my base-camp way of relating no-intended-audience cases to Grice's account, but is worse than it in two ways: (1) With few exceptions, one can see how to operate from the Gricean base-camp; whereas the corresponding task for Searle looks hopeless: how can his partial treatment of representation help us to understand the ordinary notion according to which one can represent that P as part of a deliberate attempt to deceive? (2) With my strategy, the basic central class of cases is marked off by the Gricean account, with no threat of circularity or regress. Searle's basic, central class of cases, on the other hand, contains. . . what? All Searle tells us is that in each member of the class: 'There is some P such that: S intends to be doing something that he cannot succeed in, unless P is true independently of what he is doing.' Because Searle does not say that this fits only items that fall within the class in question, his is admittedly an incomplete account of the latter. If it is even that, however, it must have some power to guide one's mind towards the class

¹ Searle, 'Meaning, Communication, and Representation', p. 215, lightly edited for brevity.

of cases that Searle has in mind; and I don't see how anyone could be thus guided by it except *through* the thought: 'Ah yes, in each of these cases the "something" that S intends to be doing is truthfully representing that P.' So I think that Searle's account, as well as being incomplete, is covertly circular.

A completely different line that Searle may want to take to deal with lying will be discussed in section 12.

11. Trouble with injunctions

Searle's account, as so far reported, does not fit injunctions. Indeed, the 'criterion of success' conditional is false of every injunction. When I say to you 'Please close the door', I presumably intend to represent your closing the door; and my success does not require that that should obtain independently of my utterance—quite the contrary. (The 'independently of S's utterance' clause is needed for statements. Without it Searle's conditional would apply whenever someone tries to make it the case that P, that is, intends something in which he must fail unless P; and that would cover so much ground other than representation as to be an intolerably feeble start on an account of representation.)

Searle knows this. When he re-runs his partial account of representation in the context of language, he leads off with the words: 'For simple indicative sentences of the sort used to make statements...' (pp. 217f), and a page later he writes: 'So far we have considered only the nature of meaning and representation as they apply to utterances of the statement class.' He proceeds to discuss four other kinds of illocutionary act; but I shall bypass three of them and concentrate on injunctions, which Searle calls "directives".

He offers a necessary condition for an utterance to be meant as a directive. If S uttered x meaning it as a directive to H to do A, Searle says, then

'In the uttering of x S intended that a criterion of success of the uttering of x will be that H does A, at least partly because of the recognition by H that S intends the uttering of x as a reason for doing A.'

This, like Grice's similar account, doesn't fit the case where S shouts 'Close the door!' without caring whether H obeys, or even (relying on H's counter-suggestibility) as a way of keeping the door open.

There are two ways in which Searle might try to deal with this.

The first way is through a base-camp strategy such as I suggested he might adopt to cover lying statements. (The present difficulty is the injunctive analogue of the liar difficulty.) Searle might say that he is offering a partial analysis of a large central class of injunctions, on the basis of which he can then explain the marginal kinds of case that do not satisfy his necessary condition—namely the ones in which the commander is not trying to be obeyed.

This is less open to criticism than the base-camp treatment of the liar, if only because where injunctions are concerned the base-camp can be established without circularity. But that very advantage also disqualifies this approach as an alternative to Grice, for it is Grice's own analysis that has established the base camp. Searle's own partial account of directives is nakedly Gricean, as indeed it has to be; I shall return to that point in section 13 below.

12. A special notion of success?

The second way of dealing with the problem is by using 'failure' and 'mistake' in explaining 'conditions of success'. Searle suggests that he means by 'success' what it ordinarily means. (See the passages quoted in section 9 above.) But perhaps he can cancel that suggestion and claim to be using 'success' as a technical term with a special meaning. If

so, then he might say that the indifferent or wily enjoiner is not a counterexample to his conditional, and that in thinking otherwise I have been construing 'success' in a naively narrow (or shallow) way. He could point out that whatever S wants H to do, if S intends to be uttering a command then he intends something such that if H doesn't close the door it 'fails' in the sense that it is *not obeyed*; and that, he could say, is a kind of failure in a command, whatever the further purposes of the commander may have been.

The problem of lying statements might also be approached in that way. When someone makes a statement (Searle might say), no matter what his ultimate purpose, the statement's being false is a kind of failure.

Hints of this line of thought can be found in what Searle writes. A little before his discussion of injunctions, he says:

The different illocutionary points can be defined at least partly in terms of the different conditions of success that the speaker intends in making the utterance. In each case the essential condition on the speech act will be specified as a necessary condition of the success of the utterance as intended, when the utterance is intended as a member of the relevant class. The basic idea here is the old one that the meaning of a statement is somehow given by its truth conditions, the meaning of a command by its obedience conditions, and so on.¹

This pretty clearly implies that if I intend my utterance as stating that P, there is some success that I cannot have unless P; and if I intend it as an injunction to do A, there is some success that I cannot have unless A is done.

On this understanding of what Searle is saying, the notion of 'success' is being divorced from that of what S is, in the ordinary sense, trying or intending to do; we have to regard his utterance as falling even if *he* is perfectly satisfied with every aspect of the situation, and not because of any relevant ignorance or error.

A newcomer to this scene could be forgiven for protesting that nothing properly called 'success' could be thus divorced from the notion of what an informed agent would count as success. We old hands know better: we have the idea of statements as fixed by truth-conditions, and of injunctions as fixed by obedience conditions, and we can think of the satisfaction of those as a kind of 'success' of the utterance, whatever the utterer may be up to. We are accustomed to thoughts like this: 'Whichever direction of fit we are interested in, a lack of fit in a particular situation is a lack of fit, and so something which is in itself unsatisfactory.'² But when we understand what Searle is saying in this way, we are going *through* our grasp of what statements and injunctions are, and so once again the procedure is circular.

Searle's book *Intentionality*, though published long before the paper I am discussing, was written long after. I have gone to the book for help with this matter, but have been disappointed. 'Conditions of success' are there usually called 'conditions of satisfaction', and 'satisfaction' must be the most frequently used technical term in the book. Searle doesn't define it, however, nor does it occur in his index. As applied to mental states, it boils down to the two theses that

- a belief is satisfied only if it is true, and
 - a desire is satisfied only if what is desired is obtained,
- and Searle makes no attempt to justify his evident assump-

¹ 'Meaning, Communication, and Representation', p. 230, quoted with harmless omissions.

² Ross Harrison, 'Ethical Consistency', in R. Harrison (ed.), *Rational Action: Studies in Philosophy and Social Science* (Cambridge University Press, 1979), pp. 29–45, p. 35.

tion that those two theses involve a single sense of the term 'satisfaction'. In defence and explanation of the peculiar notion of 'psychological state that can be satisfied' Searle offers a putative technical sense of 'satisfy' according to which 'x satisfies y' is true if y is a desire and x 'satisfies' it in the ordinary sense or *if y is a belief and x makes it true*. This seems to be incurably disjunctive; I can find no unitary sense of 'satisfy' that produces this result.

Searle wants 'satisfaction' so as to have a solid account of the genus *intentionality*, an account which neatly bifurcates down the road to yield the species belief and desire. The split, according to Searle, comes from a difference in 'direction of fit' between the state and what 'satisfies' it: with beliefs the direction runs from world to mental state, with desires it runs the other way. Having criticized the part of this story that lies upstream of the fork, I should add that the fork itself is suspect. Explaining 'direction of fit' in terms of statements and injunctions rather than beliefs and desires, Searle says that statements are 'supposed in some way to match an independently existing world' whereas imperatives are 'supposed to bring about changes in the world' (*Intentionality*, p. 7), and speaks of where the 'fault' or the 'responsibility' lies if fit is not achieved. One sees what he is getting at, of course, but these merely suggestive remarks do not locate any clean, worthwhile concept of direction of fit. From the functionalist standpoint that Searle rejects, the whole idea of a unified story that forks further downstream is mistaken. Functionalists are in a position to see (and if they bothered with details they *would* see) that the concepts of belief and desire are deeply different in structure, having almost nothing in common except the formal feature of being a psychological propositional attitude.

13. The disunity in Searle's account

Searle's partial account of meaning tells one story for statements and a different one for injunctions, with no significant overlap between the two. The word 'represent' occurs with statements and not with injunctions, and although it probably could be dragged into the latter as well, it wouldn't then answer to the only partial account Searle has given of it because that includes the 'independently of the utterance' clause.

A Gricean account respects the difference between statements and injunctions, linking one with trying to produce belief and the other with trying to produce action. But it also has an underlying unity, in the idea that each kind of attempt involves a reliance on the Gricean mechanism of intention-dependent evidence for what the speaker believes or wants. Searle could have that unity too, but he would say that it exists not on the ground floor of meaning but up at the level where communication is to be found.

It looks as though Searle cannot, down at what he takes to be the ground floor, say anything substantive about statements and injunctions at once. What I am not sure of is whether he would mind this. There is in fact a discomfort, an awkwardness, about trying to get a picture of Searle's picture of the situation. On the face of it, we are to be offered an account of illocutionary acts from which everything perlocutionary is excluded, Grice's fault being that he let the latter in on his ground floor. But the account that Searle offers as a rival to Grice's is confined to statements, and the demarcation line around statements itself involves some reference to what the utterer intends to result from what he is doing. Here is why.

The concept of an injunction cannot be explained without reference to an intended effect on an audience, that is, without reference to intended effects of the kind Searle calls

'perlocutionary'; so the concept of a statement must at least involve enough about perlocutionary intentions to ensure that statements are not injunctions. And so it is in Searle's account, where representation as it occurs in statements is partly explained in terms of a person's intending something in which he will fail unless P is the case *independently of his action*, which implies that he intends that it *not* be the case that his action or utterance make it the case that P. That intention is very abstract or general, or, as we sometimes say, it is 'negative'; but it is perlocutionary for all that. And of course when Searle moves over to injunctions the perlocutionary content gets even greater, because now the speaker has a 'positive' intention to achieve a certain perlocutionary effect.

14. Summary

Searle's use of 'representation' to explain meaning seems not to be a strong rival to Grice's because it is not and apparently cannot be accompanied by a non-circular account of what representation is. Searle does not claim to tell more than a part of the story—a necessary condition for representation—but even that is false on its most natural interpretation, and seems to be true only when understood in a manner that makes the account circular. Furthermore, the account is not even *prima facie* plausible when applied to injunctions; for them a quite different story must be told, in which "represent" has no urgent work to do. Finally, the concept of an intended audience is involved in the very distinction between statements and injunctions, and thus in the foundations of Searle's own account of meaning.¹

¹ I have been helped by the comments of William P. Alston on a draft of this paper, and by correspondence with John Searle.