

12 Austin on speech acts

Key text

J. L. Austin, *How to do Things with Words*, 2nd edn, J. O. Urmson and M. Sbisà, eds. (Oxford: Oxford University Press, 1975), especially lectures I, V, VI, VIII, and XI.

12.1 Introduction

Truth has some claim to be the central topic of philosophy. It is therefore not entirely surprising to find *philosophers* of language (as opposed to students of linguistics and grammarians, for example) concentrating particularly on truth in their treatment of language. Analytic philosophy of language may be said to begin with Frege's determination that the fundamental thing about the meaning of a sentence is its truth-value. And we've seen Davidson's related claim, that the meaning of a sentence may be given by giving its truth-conditions, forming the core of his philosophy of language.

This focus on truth has led to a corresponding focus on the kind of sentence which can be used to say something true: the *declarative* sentence – the kind of sentence which it makes grammatical sense to insert in the gap in the phrase 'Simon says that ...' It has therefore come to seem natural to regard sentences of this grammatical type as the *basic* kind of sentence, and to regard their meaningfulness as being closely connected with what is involved in their being true or false.

In a series of lectures, worked on over several years in the 1950s and eventually published as *How to Do Things with Words*, the British philosopher J.L. Austin set out to challenge this apparently natural view. He began in what may seem a peculiar way: not by focusing on sentences which are not of the grammatical type known as declarative, but by

considering sentences which seem to belong grammatically to that type, but which look odd in a more unsettling way. These sentences, which Austin called *performatives*, seem not to describe anything in the world at all, and so seem not to be true or false. Instead, they seem to *get something done*. Austin himself appeared to argue that these performative sentences cannot finally be set apart as a wholly peculiar class; but he thought that they could only be understood if we altered quite radically our conception of the nature of language.

If we focus on sentences which are capable of being true or false, we can think of sentences as things which we might hold up against the world, like pictures. This takes sentences out of the context of our everyday lives. Austin's focus on performatives – sentences which we can *do* things with – leads to a general concern with the *acts* we may perform when we use sentences, and to a whole dimension of evaluation of such acts which is distinct from the simple evaluation of statements in terms of truth and falsity. If we follow Austin, our interest in language is shifted from the concentration on truth which characterizes the bulk of work in the analytic tradition, to a general concern with the various ritual and conventional procedures involving language with which we carry on our everyday lives.

12.2 Performative utterances

Consider these sentences:

- (1) I promise that I'll be there;
- (2) I name this ship the Enterprise;
- (3) I give notice that the next meeting will be held on 1 August;
- (4) I sentence the prisoner to 14 years' hard labour;
- (5) I declare the festival open.

You will naturally expect these sentences to be spoken in particular contexts: (1) by someone giving a promise to somebody else; (2) by someone naming a ship in a public ceremony; (3) by the chairman or secretary of a society in publishing the date of a meeting; (4) by a judge as she passes sentence; and (5) by some dignitary opening a festival. These are, we might think, the *natural* uses of these sentences.

All of these sentences are of the grammatical form known as declarative. But in these natural uses they don't seem to be used to

describe the world: instead, they seem to be used to *do* something, to *perform* some action. In a natural use of (1), I don't (it seems) report the fact that I am promising: I actually promise. In a natural use of (2), I don't tell you what I call the ship: I give it a name. In a natural use of (3), I don't describe my giving notice: I actually give notice. And so on. It seems, intuitively, as if nothing true is said in these natural uses of such sentences. Instead, these uses are *performative*: they are *performances* of acts of certain kinds. Austin accepted the intuitive view that performative uses of sentences are not uses in which anything true or false is said, and so contrasted these *performative* uses with those he called *constative*. Constative uses, in Austin's sense, are precisely uses in which something true or false is said.

It's a striking fact that sentences (1)-(5) include verbs which may be used to describe acts we can perform by speaking or writing; and these verbs appear here in the first person of the present tense. But the presence of such verbs in the first person of the present tense is neither necessary nor sufficient for uses of sentences to be performative. The acts which are performed by natural uses of sentences (1)-(5) could equally well be performed by using the following sentences instead, provided that they're uttered with appropriate intentions and in the right contexts:

- (1a) I will be there;
- (2a) This ship is the Enterprise;
- (3a) The next meeting will be held on 1 August;
- (4a) The prisoner will serve 14 years hard labour;
- (5a) The festival is now open.

On the other hand, there are uses of (1)-(5) in which they can be taken to be reports of fact. Imagine an entry in the diary of a busy judge and local celebrity:

10 am: call from festival organizer worried about the afternoon ceremony; I promise that I'll be there. 10.30: rush to shipyard, where I am hustled onto a platform beside a huge ship; as rather boringly instructed, I name this ship the Enterprise. 12 noon: brief preliminary meeting of GP executive committee; I give notice that the next meeting will be held on 1 August. 2 pm: in court for sentencing in Abercrombie case; I sentence the prisoner to 14 years hard labour. 4.35: after a quick change, I arrive at the festival site, to be greeted by the organizer (extremely worried, because I'm 5 minutes late); I declare the festival open. 7 pm: back home with a splitting headache; and so to bed.

These rather simple-seeming points raise a number of quite large issues about language. The first and most straightforward one is the nature of the distinction between performative and constative uses of sentences. There seems a very clear difference here: the use of sentences (1)–(5) in the imagined diary entry seems clearly different from their natural, performative, uses; and sentences (1a)–(5a) all have quite simply descriptive uses, which seem clearly distinct from the performative uses in which they seem almost equivalent to sentences (1)–(5). But how precisely is that distinction to be characterized? And here's a related question: is Austin right to think that performative uses do not say anything true or false? If I say, 'I promise', do I merely promise, or do I also say truly *that* I promise?

There are larger questions about meaning involved here too. Crucially, consideration of performatives shows that there is a difference between the meaning of words and sentences, on the one hand, and the meaning or significance of *uses* of words and sentences, on the other. It seems that there's no ambiguity of meaning in the word 'promise' between the performative and the diary-entry uses of (1), or between its use in (1) and its use in either of these two sentences:

- (1b) I promised that I would be there;
- (1c) She promises that she'll be there.

And yet there seems a clear difference in what's being done between the performative and diary-entry *uses* of (1), or between the performative and the obviously descriptive *uses* of (1a)–(5a). What, then, is the relation between the significance of a *use* and the meaning of a word or sentence? Does the use presuppose the meaning of the words, or is the meaning of words to be explained in terms of the significance of uses of them?

12.3 Towards a general theory of speech acts

Austin began with the idea that there was an important contrast between *performative* utterances – in which something is *done* in the uttering of words – and *constative* utterances – which can be true or false. Clearly this distinction is a conflation of two different distinctions, which might, in principle, diverge. On the one hand, there's the supposed distinction

between utterances in which something is done and those in which nothing is done (or nothing like that). And, on the other hand, there's the supposed distinction between utterances of declarative sentences which are true or false and utterances of declarative sentences which are not true or false. In this section, we'll look further at the performative/non-performative distinction. In the next, we'll consider the claim that performative utterances cannot be true or false.

Despite beginning with the idea of a contrast between performative and constative utterances, Austin ended up concluding that there was no way to characterize performative utterances which did not count constative utterances as performative.¹ What does a performative utterance have to be like? We might think that it needs a special performative *verb* – a verb which describes a kind of act which might be performed by speaking, such as 'promise', 'order', 'baptise', and so forth. And we would then expect that verb to be in the first person of the present tense. But we've already seen that this condition is neither necessary nor sufficient for an utterance's being performative. Not necessary, because an utterance can be performative without containing such a verb in the first person of the present tense: (1a)–(5a) can be used in performative utterances, for example (and there are also performatives which use the second or third person – 'You are hereby warned', for instance – as Austin points out).² And not sufficient, because utterances of sentences which meet these grammatical criteria need not be performative, as the example of the judge's diary shows.

Austin's eventual criterion was something like the following. Utterances of sentences which contain no performative verb – sentences like (1a)–(5a) – are performative if they are 'equivalent', in some sense, to utterances of sentences which do contain a performative verb – sentences like (1)–(5). Take the first in each list of examples, in particular, beginning with the sentence without a performative verb:

(1a) I will be there.

¹ F. Recanati thinks that Austin's change of mind depends, in effect, on a change in the meaning of 'performative'; see his *Meaning and Force: The Pragmatics of Performative Utterances* (Cambridge: Cambridge University Press, 1981), pp. 70–2.

² *How to Do Things with Words*, p. 57.

This sentence might be used to make a prediction ('I will be there, I expect – I usually go'). But if it's used performatively, that use will be equivalent, in some sense, to a use of the following sentence:

(1) I promise that I'll be there.

Again, (1) could be used non-performatively (as in our judge's diary entry), but not in a use which is equivalent to a use of (1a). And we can make (1) unambiguously performative by means of some self-referential device, as in the following sentence:

(1*) I *hereby* promise that I'll be there,

even if this might seem rather formal for everyday use.

Austin's idea is that a sentence like (1a) can be used in what he calls a *primary* performative (the kind of performative utterance which could be made before there were words to describe types of linguistic act). The performative is then made *explicit* in the reformulation involving a performative verb (like (1)). The thought is that in an explicit performative someone performs the same act – promising to be there, for example – as is performed in the corresponding primary performative: it is just that the words make it explicit what act is being performed.

If this is the test of performativity, then it seems that Austin's original 'constative' (statement-making or descriptive) utterances will count as performative too. Consider the following sentence:

(6a) The cat is dead.

A statement-making utterance of (6a) is a constative utterance, if anything is. But if 'promise' is a performative verb, because it describes an act which can be performed in speaking, then surely 'state' is also a performative verb. And in that case, it seems that I can make explicit what I'm doing in uttering (6a) by uttering the following sentence instead:

(6) I state that the cat is dead.

Austin would regard the use of (6) as 'equivalent' to the use of (6a) in just the same sense as the use of (1) is 'equivalent' to the use of (1a) to make a promise. So it seems that constatives are performatives too.

Austin's reaction to this is to move from the rather specific contrast he began with, between performative and constative utterances of declarative

sentences, to the outline of a general theory of speech acts. In any ordinary use of language, he suggests, a speaker will be performing acts of at least two, and possibly three, importantly different kinds.

In the first place, she will be performing what Austin calls a *locutionary* act. This is more than merely uttering sounds: it's speaking the words with the meaning they have (or have here). It's what we might call *saying something*. Secondly, as Austin puts it, she will be doing something *in* saying that. In using those words, she might be asking a question, giving an order, making a promise, stating a fact, and so on. These are all what Austin calls *illocutionary* acts. And finally, she might achieve something *by means of* saying what she says: she might draw someone's attention to something, convince her of something, get her to do something, and so on. Austin calls these acts of achieving something by means of saying something *perlocutionary* acts.

When someone speaks she may be performing acts of all of these kinds. Suppose someone utters the sentence 'Shut the door!' If she uses this sentence with the meaning it standardly has, she has, as we might say, *said something*: she's performed a *locutionary* act, rather than merely made some noises. In all probability, she's also given an order: this is an *illocutionary* act. And it may be that the order is obeyed, and her audience shuts the door. In that case, she's performed the *perlocutionary* act of getting someone to shut the door. There's a sense in which these are all different acts, in that they're all different things which she can be said to have done. But she doesn't need to do them all separately: she does all these things just by uttering the sentence 'Shut the door!'

Austin's focus on the acts which people perform when they speak can seem like a breath of fresh air in the intense atmosphere of analytic philosophy of language. We seem to have moved away from the logic-oriented, theoretical approach to language which has dominated the analytic tradition; instead language seems to be placed in the middle of real lives. This reaction is understandable, but it needs to be treated with some caution. Austin was, in fact, concerned to distance himself from *two* different approaches to language, not one. One part of the analytic tradition has concentrated simply on the meanings of words and sentences: this is the line which descends from Frege, whose concern is with semantics. But another part, reaching back to Locke at least, has always been concerned with the role of language in people's lives. Language has standardly been thought to have a function: immediately

one of communication, but ultimately one of making people's lives safer and better. Too strong an emphasis on this idea of the *ultimate* point of language is itself one of the things which Austin opposes.

Austin's particular concern is with the class of *illocutionary* acts – the questionings, orderings, promising, and statings which may be performed in uttering sentences. He thinks there's a tendency to try to assimilate illocutionary acts to acts of one of the other two kinds – either to mere sayings, or to the things we achieve by means of saying things. The tendency to assimilate illocutionary acts to locutionary acts – mere sayings – may be associated with the theoretical, semantically-oriented line of the analytic tradition. Austin himself is not at all sceptical about this tradition. Indeed, he seems to adopt without question the idea that words have both Sense and reference, and is in that respect an orthodox Fregean. His point is rather that concentration on the meaning of words misses something central about language.

Nor is that central thing to be understood by thinking of language as a tool by means of which certain desirable results can be achieved. To rush too quickly from concentration on locutionary acts – mere sayings – to focus on what can be achieved by means of language is to ignore the importance of illocutionary acts from the other side, on Austin's view. This is manifested in the tendency to assimilate illocutionary acts to *perlocutionary* acts. Austin's concern is to make us focus on *illocutionary* acts, which are, in a sense, intermediate between the locutionary and the perlocutionary.

But why should the illocutionary act be particularly important? Consider the two obvious uses of a sentence like (1a):

(1a) I will be there.

We expect this to be uttered in the making of a promise, but it could (given a suitable context) be used to make a prediction. On Austin's view, if an utterance of (1a) is the making of a prediction, then we will have something which is true or false (according to whether or not the utterer is, in fact, in the place in question at the relevant time); but if it is the making of a promise, there will be nothing for which the question of truth and falsity even arises. Whether or not we have a bit of language (to put it deliberately loosely) which is capable of truth or falsehood seems to depend on which illocutionary act is being performed. If this is right, it has both a specific significance in the history of the philosophy of

language, and a general significance for understanding what matters about language.

The specific significance is that if Austin is right, it will at least be more complicated to take a Fregean or Davidsonian approach to semantics, and explain the meaning of words in terms of the truth of sentences. For declarative sentences, on their own, will not be true or false: we will only have something capable of truth or falsity once an appropriate illocutionary act has been performed.

The general point about language follows on from that. If Austin is right, many of the most important features of the language we encounter will depend on the illocutionary act being performed, rather than on the meaning of words and sentences. When we're faced with someone saying something, it's evidently of fundamental importance whether we're being faced with a question or a statement or a command. The issue will not be determined just by grammar – though we might expect grammar often to provide a clue – but will depend on what illocutionary act is being performed.

12.4 Truth and performatives

On Austin's view, both of the following sentences may be used to make the same promise:

- (1) I promise that I'll be there;
- (1a) I will be there.

And both of the following sentences can be used to make the same statement:

- (6) I state that the cat is dead;
- (6a) The cat is dead.

According to Austin, nothing true is said by either (1) or (1a) (on this use) and just one thing is stated by both (6) and (6a) – that the cat is dead – and this can be true or false. The role of the phrases 'I promise that' and 'I state that' is, on his view, just to make explicit what illocutionary act is being performed in the utterance of the sentence which follows. If I utter (1) in the way which makes it roughly 'equivalent' to (1a), or (6) in the way which makes it roughly

'equivalent' to (6a), I do not state *that* I promise or *that* I state, nor do I describe myself as promising or stating.

The difficulty with this view is in finding an answer to Davidson's famous question (asked in another context): what are these familiar words doing here?³ How can we explain what the word 'promise' is doing in (1) or what the word 'state' is doing in (2), while both respecting Austin's view and avoiding treating these words as being ambiguous? It's clear that these words do *something*: at the very least they introduce intensional contexts in (1) and (6) which are not present in (1a) and (6a). It's not obvious what should be said. In the light of this difficulty, some have supposed that utterances of (1) and (6) do, after all, make statements about the illocutionary acts being performed.

Kent Bach and Michael Harnish, for example, suggest that in uttering (1) I state that I promise that I will come.⁴ This statement is true just in case I do, in fact, promise that I will come. Where and when do I actually make the promise? We may suppose that (1) is the only relevant sentence I utter. In that case, the promise must be contained in (1) itself. That is to say, in uttering (1) I make a statement (that I promise) which is made true by my uttering of that very sentence, (1), itself. We can now explain what the word 'promise' is doing in (1): it is doing exactly the same as the thing it does in the following two sentences:

(1b) I promised that I would be there;

(1c) She promises that she'll be there.

According to Bach and Harnish, in uttering (1) to make a promise, I'm really performing *two* illocutionary acts: I'm stating that I promise to be there, and – because that statement is true – I'm promising to be there. The same analysis is applied to utterances of such sentences as (6).

I think the principal reason why we might feel uncertain about this is probably what underlies Austin's insistence that if I utter (1) to make a promise, I don't state anything. (Austin himself offers no clear reasons.) Bach's and Harnish's view requires such an utterance of (1) to be *self-verifying*. There's a common intuition that what makes a statement true

³ D. Davidson, 'On Saying That', *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984), p. 94.

⁴ K. Bach and R. M. Harnish, *Linguistic Communication and Speech Acts* (Cambridge, MA: MIT Press, 1979), pp. 203–9.

must be somehow independent of that statement itself. This is part of what lies behind the correspondence theory of truth, according to which a true statement is one which *corresponds* to a fact.⁵ It seems part of the notion of correspondence that correspondence involves two independently existing entities, which happen to be correlated, but might not have been. Bach's and Harnish's suggestion seems to conflict with this correspondence intuition.

The intuition is not indubitable, though it's not indefensible either. The issues here are complex, and I can only hint at some of them here. We might note that if anyone manages to mean anything by uttering 'I exist', it seems bound to be true, and this might be thought to be a form of self-verification. But this doesn't seem to violate the original intuition, because the point remains that a person's existence is independent of her *saying* she exists. Again, it's natural to think that if someone says 'I am speaking', her utterance is self-verifying. But in most actual cases in which we might imagine this sentence being used (by an exasperated teacher addressing a class, for example), it's the utterance of the *surrounding* sentences, rather than that one, which is naturally understood to make the statement true. On the other side, we might think that there's a connection between this issue and the liar paradox. The liar paradox arises in the case of sentences like this:

(L) (L) is false.

It's not unnatural to think that the liar paradox is to be solved by insisting, in the manner of the correspondence theory, that what makes a statement true must be independent of that statement itself. But the proper treatment of the liar paradox is an enormous subject on its own; at the very least, it's far from obvious that there's no way of solving it while accepting Bach's and Harnish's account of performatives.⁶

Is there any alternative to Bach's and Harnish's account of what the word 'promise' is doing in (1) or the word 'state' in (6), which might allow us to follow Austin in denying that a promising use of (1) is true or false, or that an explicitly performative use of (6) states more than one thing? One suggestion that is worth pursuing is introducing the idea of special speech

⁵ Austin himself seems to have accepted some form of correspondence theory of truth: see his 'Truth', *Aristotelian Society Supplementary Volume*, 24 (1950), pp. 111–29.

⁶ For a brief introduction to the liar paradox, see R. M. Sainsbury, *Paradoxes*, 2nd edn (Cambridge: Cambridge University Press, 1995), ch 5.

acts of *referring*. We might suppose that whatever other speech acts we perform when we speak – acts of promising, stating, questioning, or whatever – we always also perform acts of referring by means of our words. We might here adopt a relatively orthodox world-directed view (in the anti-Fregean tradition started by Russell), and take uses of singular terms to involve acts of referring to objects, uses of predicates to involve acts of referring to qualities and relations, and uses of sentences to involve acts of referring to combinations of objects and qualities or relations, which we might call *states of affairs* or *situations*.⁷

This kind of theory faces two sorts of difficulty if it is to be offered as an alternative to the Bach–Harnish view. First, it needs to be explained how sentences can refer to states of affairs without being true or false. And secondly, it needs to deal with the various difficulties which have led Davidson and others to abandon referential accounts of meaning.⁸ But if these problems can be dealt with, it may be that Austin’s original intuition can be preserved.

12.5 Issues for a theory of speech acts

Austin’s emphasis on the importance of illocutionary acts raises some questions for everyone concerned with the philosophy of language, and some questions specifically for those who aim to develop a theory of speech acts.

One central issue arises over the meaning of words. It seems undeniable that the meaning of words is, in some sense, prior to each individual illocutionary act. When I perform an illocutionary act, I exploit this prior meaning of the words I use. Austin himself seems to have a simple, conservative view of the meaning of words and sentences: they just have Sense and reference. He seems to suppose that this is something quite independent of speech-act theory. The difficulty is that it seems that truth and falsity only enter the picture once it has been determined that an illocutionary act of stating is being performed. This means that it’s not

⁷ A view of this general kind is offered by Stephen Barker, *Renewing Meaning: A Speech-Act Theoretic Approach* (Oxford: Oxford University Press, 2004).

⁸ For Davidson’s worries, see the opening pages of ‘Truth and Meaning’, in his *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984), pp. 17–36. A related set of worries was considered in ch. 9, § 9.2.

clear how the meaning of words can be explained in terms of the truth of sentences. Indeed, it's hard to see how sentences can be true at all, on this view: they cannot even be true relative to an *occasion* of utterance, because it's the kind of illocutionary act being performed, rather than merely the reference given to each of the component words, which determines whether we have something capable of truth and falsity.

This raises a difficulty for everyone – including, most obviously, Frege and Davidson – who hopes to explain the meaning of words in terms of the truth of sentences. What else might we propose? There are two natural alternatives.

First, we might adopt some form of referential theory of meaning, as has just been suggested in offering an alternative to the Bach–Harnish view. We might suppose that singular terms refer to objects, predicates to properties and relations, and sentences to states of affairs or situations. Again, it will have to be claimed that we can make sense of sentences referring to states of affairs without being true or false; and we will have to deal with the traditional worries about referential theories.

Or, secondly, we might try building speech-act theory more thoroughly into our conception of semantics. We might attempt to explain the meaning of sentences in terms of the illocutionary acts that *could* be performed with them. Meaning will be explained in terms of illocutionary-act *potential*.⁹ This suggestion itself is not without difficulty, however. After all, we've seen that the very same sentence can be used to perform quite different illocutionary acts – even when the words seem to have the same meaning. This was what we found with the groups of sentences (1)–(5) and (1a)–(5a).

Here's a natural way of dealing with that difficulty. Consider this sentence again:

(1a) I will be there.

If I use this to make a promise, the promise I make will be the promise *that I'll be there*; and if I use it to make a statement, the statement I make will be the statement *that I'll be there*. This idea can be generalized. If I use a sentence to ask a question, then *what* question I ask will be determined by

⁹ This approach is championed by William Alston in his *Illocutionary Acts and Sentence Meaning* (Ithaca: Cornell University Press, 2000), although the details presented below are not his.

the meaning of the sentence. If I use a sentence to give an order, then *what* order I give will be determined by the meaning of the sentence.

We may then suggest that the meaning of a sentence is a matter of what *particular* illocutionary act would be performed by it, once it is determined what general *type* of illocutionary act it is. Can we perhaps formalize that suggestion, with a view to giving an explicit semantic theory for a language? Here's how we might begin, at the level of sentences at least.

(IAM) A sentence *s* means that *p* if and only if

For any illocutionary act type *A*, anyone who *As* in uttering *s* thereby *As* that *p*.

So it will follow that, if *s* means that *p*, if someone makes a *promise* in uttering *s*, she will be promising that *p*, and if someone makes a *prediction* in uttering *s*, she will be predicting that *p* – and so on. If we begin with something like (IAM) as the account of the meaning of sentences, some work will have to be done to explain what it is for words to have meaning, and how the meaning of sentences is dependent on the meaning of the word of which they are composed. But it may be that existing semantic theories can be adapted to the purpose.¹⁰

If this and some form of referential theory are the two natural options for a speech-act theorist who's concerned to explain word-meaning, it's arguable that speech-act theory introduces nothing very radically new to the study of semantics.¹¹ The general shape of the options seems not much different from what was available before we considered speech acts: on the one hand, a referential theory; on the other, a theory which aims to use sentences to state the meaning of sentences within statements of the familiar form '*s* means that *p*'.

In that case, the fundamental task for speech-act theories is to understand the character of the acts which are performed in the use of words and sentences. One central question is this: what determines what illocutionary act is performed when someone speaks? It's natural to think that the intention of the speaker has something to do with it. A further

¹⁰ Note that an extensional theory, like a Tarskian truth-theory will face at least the usual problems, since '*x As that ...*', where '*A*' is schematic for some illocutionary verb, introduces an intensional context.

¹¹ Though presumably this claim would be challenged by Barker: see his *Renewing Meaning*.

question is whether an illocutionary act needs to conform to rules established by convention. Austin's view seems to have been that the various kinds of illocutionary act are established by convention, and someone counts as performing a particular kind of illocutionary act (promising, say, or asserting) in virtue of meeting the conventionally established conditions for performing acts of that kind. It is not entirely clear what kind of convention Austin had in mind, but we can develop something which seems in the spirit of his work by beginning with the case of promising, and generalizing some of its crucial features.

At the beginning of his book, much of Austin's focus is on certain kinds of illocutionary act which seem obviously institutional. The naming of a ship, the sentencing of a criminal, the official opening of a festival, giving notice of a meeting, are all acts which take place within institutions. Acts of these kinds are evidently bound by rules, which set boundaries for the acts' being properly performed: they need to be performed by the right people, in the right way, on the right occasions, and so on. The institutions and the rules which bind them are naturally thought to be conventional. These institutions are not natural objects – they don't simply grow like trees – so their constitutive rules are, in a sense, arbitrary: other institutions, with different constitutive rules, could have been established. Moreover, these institutions are established and kept in place by the agreement and connivance of their members and the people who interact with them.

Promises are not exactly like these obviously institutional acts, but they may seem to have some affinity with them. We may speak of an institution of promising – the fact that we give and accept promises, and act on the basis of them – but this isn't an institution like a club or society, nor is it one like a judicial system. Nevertheless, the practice of making promises need not have existed, and seems to depend for its continued existence on the attitudes and behaviour of the group of people who accept each other's promises. Within this 'institution' we can recognize certain kinds of rules for promises to be genuine promises: we can see that there are questions about who can legitimately promise what, and on whose behalf, for example.

It's not obvious that the same applies to the large majority of illocutionary acts, however. It's natural to think, for example, that the illocutionary acts of informing and of asking questions are essential to the nature of language itself. There may be some sense in which language as a

whole is conventional, but that doesn't immediately make these particular types of illocutionary act conventional. Austin's view that all types of illocutionary act are conventional seems to require that there could be languages without acts of informing or questioning: these have to be thought of as particular routines which happen to have sprung up in particular languages and are sustained by the consent of their speakers – although they could, in principle, be abandoned.

Others oppose Austin's view. According to Bach and Harnish, for example, the fundamental types of illocutionary act are fixed by the states of mind which they express; and there need be nothing conventional about them. One fundamental type (Bach and Harnish call them *constative*, for obvious reasons) are expressions of *belief*: if I assert that the cat is dead, for example, I express the belief that the cat is dead. Bach and Harnish call another fundamental type *directives*, which they define as expressions of the speaker's attitude towards some prospective action by the hearer:¹² if I ask Winnie to shut the door, for example, I'm expressing my desire that she shut the door. On their view there are no specific procedures which have to be followed to perform an act of one of these types. There are no conventional rules which have to be followed in order to perform them.

So what is it to perform an illocutionary act, on such a view? It is, fundamentally, to intend to do something which is an expression of a particular attitude (belief, desire, or whatever). If we put this together with the preliminary account of sentence meaning proposed in (IAM), we have an account of meaning on something like the following lines. What a sentence means is a matter of what would be believed, desired (or whatever) by someone who really had the attitude she intended to express in uttering the sentence. This anticipates the account of meaning proposed by H. P. Grice, which is the subject of the next chapter.¹³

Further reading

The classic text in speech-act theory, after Austin's *How to Do Things with Words*, is John Searle's *Speech Acts* (Cambridge: Cambridge University Press,

¹² Their full taxonomy of types of illocutionary acts is to be found in *Linguistic Communication and Speech Acts*, ch. 3.

¹³ This is hardly surprising, of course: Bach and Harnish make extensive use of work on the Gricean account of meaning.

1969). Searle follows Austin in thinking that speech acts presuppose institutions with rules for performing particular types of illocutionary act. This view is opposed by P.F. Strawson, 'Intention and Convention in Speech Acts', in his *Logico-Linguistic Papers* (London: Methuen, 1971), an article which is also significant in the development of Grice's account of meaning (the topic of chapter 13 below). Kent Bach and Michael Harnish develop a theory of speech acts which follows Strawson rather than Searle in this respect: K. Bach and R.M. Harnish, *Linguistic Communication and Speech Acts* (Cambridge, MA: MIT Press, 1979). A recent revival of the Austin-Searle tradition in this respect is W. Alston, *Illocutionary Acts and Sentence Meaning* (Ithaca: Cornell University Press, 2000): this work also develops an 'illocutionary act potential' theory of sentence meaning. A thoroughly worked (and therefore quite technical) attempt to do speech-act semantics, using the idea that sentences represent states of affairs, is S. Barker, *Renewing Meaning: A Speech-Act Theoretic Approach* (Oxford: Oxford University Press, 2004).