

Introduction

0.1 Motivation

Anyone who has come into even the most superficial contact with pragmatics will have heard of speech acts. Language, we learn as undergrads, does not only represent the world; it also allows us to do things: to conjecture and to affirm, to command and to supplicate, to promise and to threaten, to baptise and to make oaths – to perform speech acts. Every good introductory textbook on pragmatics contains a chapter on speech acts, and the names of Austin and Searle, the founding fathers of contemporary Speech Act Theory, are often among the first ones we learn as naïve students in linguistics or philosophy of language. Beyond pragmatics, the notion of speech acts is used in syntax and semantics, in literature and cinema studies, in ethics and epistemology, in clinical and experimental psychology, and the list could be continued for a long time.

However, in spite of the acknowledged importance of speech acts, their study occupies quite a strange place within the contemporary theoretical landscape. Over the past thirty years pragmatics has seen important methodological and conceptual changes. Growing attention has been paid to the cognitive aspects of utterance interpretation, and many exciting areas of research and new empirical phenomena have been uncovered. Keeping to the general development of cognitive science, scholars in pragmatics aim at building psychologically plausible models where any component of the meaning accessed by the hearer is analysed as the output of a cognitive process that ought to be explained in rigorous, naturalistic terms. There is no reason why speech acts should be left out of this new research programme.

Studying speech acts is perhaps too often seen as an attempt to unveil natural classes. But one should not forget that, fundamentally, the object of inquiry is an aspect of the interpretation of communicative stimuli. As a purely terminological safeguard, it is convenient to equate the speech act performed by way of an utterance with the *illocutionary* force of this utterance: that is, with the illocutionary force this utterance is interpreted as having. (This way of speaking originates from Austin himself. Accordingly,

2 Introduction

the terms *speech act* and *illocutionary act* are often used synonymously, and I will follow this useful convention, too.) Now, comprehending an utterance as having a certain illocutionary force forms an essential part of communication. For instance, when addressed with *You leave tomorrow morning* one needs to know whether it is an assertion, an order or a promise. Pragmatic theory ought to provide a clear view of the cognitive mechanisms that come into play when an utterance is interpreted as a certain kind of speech act, and of the linguistic and extra-linguistic factors that determine such an interpretation. Yet, most cognitively oriented pragmatic models devote hardly any space to the way utterances are assigned illocutionary force (see, e.g., Sperber and Wilson, 1995; Levinson, 2000; Carston, 2002; Recanati, 2004, 2010; Jaszczolt, 2005). This lack of concern is surprising – all the more because most of the current discussions about the semantics/pragmatics interface presuppose some distinction between semantic content and speech act content. The ambition of this book is to fill the gap, and to provide a linguistically, philosophically and psychologically viable account of illocutionary force attribution.

0.2 Conventional and non-conventional speech acts

Some may argue that if speech acts fall outside the subject matter of cognitive pragmatics it is for a good reason. Cognitively oriented theories typically seek cross-cultural explanations, which do not depend on any group-specific convention. Now, it is well known that in the series of lectures that gave the impetus to the contemporary study of speech acts, Austin (1975) defined them as being distinctively conventional. However, equally famous is Strawson's (1964) plea for a distinction between conventional (or institutional) and non-conventional (or non-institutional) speech acts. In a nutshell, his criticism of Austin runs as follows. Understanding that an utterance amounts to a conventional speech act – for example, bidding five no trumps or opening a meeting – requires knowing that certain conventions, peculiar to a certain group, are in force. By contrast, in order to recognise a non-conventional illocutionary act, it is sufficient, according to Strawson, to grasp a certain multi-layered Gricean communicative intention.

I will return to Gricean intentions in a moment, but let me say a few words about conventional speech acts in order to make clear what this book will *not* address. In order to be performed with success, conventional illocutionary acts require a system of 'rule- or convention-governed practices and procedures of which they essentially form parts' (Strawson, 1964: 457). For instance, the utterance of *I baptise you John* does not count as baptising when uttered by someone other than the priest and/or on some other occasion than the ritual of baptism; this utterance – the speech act of baptising – has a fixed

and essential part to play within the frame of that ritual. Conventional speech acts consist in complex actions that are governed by extra-linguistic, group-specific conventions which happen to require the performance of a certain linguistic utterance (see Urmson, 1977; Bach and Harnish, 1979: 108–16; Alston, 2000: 89–95). Although the way conventional speech acts emerge and persist is an intriguing question, in this book I will exclusively address non-conventional speech acts. That is, only those illocutionary forces whose attribution can be analysed in cross-cultural terms will occupy me. In what follows, the terms *illocutionary force* and *speech acts* are implicitly restricted to the non-conventional side of Strawson's divide.

In fact, I will focus on only three main groups of non-conventional speech acts: constatives, directives and commissives. The first two are the most uncontroversial illocutionary categories and are clearly non-conventional. By *constatives* I mean those speech acts that are aimed at information transmission: assertions, testimonies, conjectures, guesses and so on.¹ As for *directive* speech acts, most of the time (we will see that the qualification is important) they are attempts at provoking an action from the addressee (A): orders, requests, pleas and so on. Finally, *commissive* speech acts commit the speaker (S) to a certain course of action: the paradigmatic examples here are promises and threats. Analysing commissives as being non-conventional is more controversial; however, I will argue that there are compelling reasons for doing so.

I take these three categories as given, without engaging in any kind of preliminary classificatory enterprise. My objective is to account, in precise and well-defined terms, for the way these illocutionary forces are attributed to utterances. In the course of the analysis, these three categories will naturally emerge as three different aspects of our communicative competence. I will not attempt to draw fine distinctions within each of these three classes: but I will try to show that the theory to be defended below is powerful enough to include every relevant subcase without over-generating.²

Two major kinds of non-conventional speech acts will be left aside: questions and exclamations. The only reason – or rather excuse – for this limitation is a practical one. It is ambitious and difficult enough to try to build an analytical framework sufficiently robust to account for constatives, directives and commissives. I hope, however, that the resulting theory can be extended to other kinds of speech acts without altering its essential features.

¹ Other authors, such as Searle (1975) and Alston (2000), name such speech acts 'assertive'. Like Bach and Harnish (1979), I prefer the term *constative*, because it does not misleadingly suggest that the class is limited to 'strong', assertion-like speech acts.

² Finer-grained speech act classifications actually presuppose a number of theoretical commitments and analytical objectives (Kissine, forthcoming, for a survey), so that one may be sceptical about taxonomic studies of speech acts (e.g. Jaszczolt, 2002: 304–9).

0.3 The two main ingredients of the account

The division between conventional and non-conventional illocutionary acts is not always as straightforward in the literature as it is in Strawson's view. Take Searle's (1969, 1995) famous definition of convention: Y is a conventional fact if, and only if, Y occurs as the right member of the formula 'In context C, X counts as Y'. Now, any illocutionary act is an utterance that is interpreted in a certain way in a certain context. Under such a conception, is it not the case that any illocutionary act is conventional? In Searle's view, it is, all the more because for him the illocutionary force is built into the sentence meaning by linguistic conventions. That is, for Searle, any well-formed sentence-type corresponds to a speech-act type in virtue of linguistic conventions (for a detailed discussion, see Recanati, 1987: 219–24; 2001a, 2003).

Linking illocutionary force to linguistic form is precisely what I will not do in the pages that follow (for supplementary criticisms of this position, see Kissine, 2011a, 2012). One of the two leading ideas of this book is that one should distinguish between the linguistic meaning of an utterance, the *locutionary* act this utterance constitutes, the illocutionary force it receives, and its potential *perlocutionary* effects. It turns out that some utterances that are endowed with contextually determined, fully propositional meaning do not constitute direct illocutionary acts; at a literal level some are just locutionary acts. As a consequence, I will argue that illocutionary forces are optional pragmatic components of the process of interpretation, which, however, are automatically triggered under certain conditions.

I have mentioned above that Strawson defended a Gricean analysis of non-conventional speech acts. Following his lead, illocutionary force attribution is often reconstructed as a more or less trivial extension of Grice's (1957) analysis of non-natural meaning, based on the attribution of multi-layered communicative intentions. To date, the most articulated version of this view, understood as a psychologically valid theory of utterance interpretation, appears in Bach and Harnish's (1979) book. According to these authors, an utterance *u* counts as expressing the mental state $\Psi(p)$ if, and only if, this utterance is produced with the intention that the addressee (A) takes *u* as a reason to believe that the speaker (S) entertains $\Psi(p)$, and that A does so because he recognises that *u* was performed with this very intention. This Gricean notion of expression serves, in Bach and Harnish's theory, as a classificatory principle; once A has identified the mental state S expressed by her utterance, he knows which type of speech act S performed. Constatives express S's intentions to make A believe that *p* is true and/or S's belief that *p* is true; directives express S's desire, intention or wish that A bring about the truth of *p*; commissives express S's intention to bring about the truth of *p* (and the belief that her utterance commits her to do so).

In what follows we will see again and again that such Gricean theories are widely implausible given our current understanding of the ontogenesis of pragmatics and of mindreading. The conditions under which an utterance is interpreted as a constative, a directive or a commissive illocutionary act – and this is the second main ingredient of my account – depend on whether this utterance counts as a reason to believe or a reason to act with respect to the conversational ground. Such an approach implies that two cognitive capacities are required in order to comprehend utterances as speech acts: (a) an ability to conceive a background of alternative models of reality; (b) an ability to conceive such models with respect to shared information. An important feature of this theory is that no attribution of complex Gricean intentions is required to grasp illocutionary force.

0.4 Outline of the book

The aim of the first chapter is to use Donald Davidson's theory of action to redefine J. L. Austin's notions of *locutionary* and *perlocutionary* acts, and to lay the basis for the theory of illocutionary force attribution to be developed later on. Analysing Austin's distinctions through Davidsonian lenses makes it clear that definitions of illocutionary acts should not encroach on locutionary or perlocutionary territory. On the one hand, some utterances have a locutionary status without having any direct illocutionary force; on the other hand, we can form intentions to perform an illocutionary act without *eo ipso* intending to produce the perlocutionary effect that is typically associated with this illocutionary act type.

Locutionary acts, I will argue towards the end of the first chapter, must fulfil two requirements: they are constitutive of illocutionary acts and they have a full-blown propositional content. Although locutionary acts have much in common with what could count as the expression of an Intentional³ state (e.g. a belief, a desire or an intention), it is possible to perform a locutionary act without expressing any such state (at least on any intuitive understanding of expression). In Chapter 2, I argue that locutionary acts are functionally isomorphic with Intentional states. Both locutionary acts and Intentional states can be functionally characterised as presenting a propositional content under a certain mode. While traditional typologies of Intentional states have been cast in terms of direction of fit, I will argue that Intentional states are better divided into potential – those, such as desires, that are neither true nor

³ Following Searle's (1983) useful convention, I will use *Intentionality* and *Intentional* with an upper case to refer to the property of mental states of being 'about facts' of the external world. For instance, beliefs, desires and intentions are Intentional states, whereas joy and anguish can be non-Intentional: viz, devoid of any object.

6 Introduction

false with respect to the subject's beliefs – and non-potential – beliefs and intentions. Locutionary acts can be divided in the same way: utterances whose function is to provide information about the world, and those whose function is to present a state of affairs as being merely potential. We will also see that some sentence-types – for instance, grammatically imperative ones – are suited to the expression of potential contents only. In this way, it is possible to understand how an utterance can constitute a locutionary act without having direct and literal illocutionary force.

The third chapter is devoted to constative speech acts. Once the locutionary and perlocutionary levels have been severed from the illocutionary level, there is not much theoretical space left in order to describe the core common to this 'information-sharing' family of illocutionary acts. On the one hand, the definition must not make reference to perlocutionary effects; on the other hand, simply equating constative speech acts with a representation of a state of affairs or an expressing of a belief would make the definition way too liberal. The account I will develop owes much to Stalnaker's (1978) classic theory of assertion. I fully endorse Stalnaker's idea that speech acts can be understood only with respect to a conversational background, conceived as the set of possibilities (of possible worlds) that remain open relative to what is mutually accepted as true in the context of the conversation. However, I reject his assumption that assertions should be thought of as attempts to modify this conversational background. The main claim of Chapter 3 is that assertions, and constative speech acts in general, are reasons to believe, with respect to the conversational background, that their content is true. This definition correctly predicts that some utterances of grammatically declarative sentences are devoid of any direct and literal constative illocutionary force; these are just non-potential locutionary acts.

Chapter 4 extends this approach to directive speech acts, which are defined as reasons for A to act (again, relative to the conversational background). An important advantage of this definition is that it dissociates features of sentence-types, such as the imperative mood, that are prototypically associated with directive speech acts from the directive illocutionary force. It correctly predicts that some grammatically imperative utterances do not involve the performance of any directive speech act; they are interpreted as mere potential locutionary acts.

As I said above, the ambition of this book is to build a psychologically valid theory: in Chapter 5 I provide empirical support to my account of illocutionary force attribution. This account implies that in order to interpret utterances as constative or directive speech acts, one needs to adopt a shared perspective on the world. We will see that in the light of the available data it is reasonable to think that this cognitive capacity is dysfunctional in persons with autism spectrum disorders, while it is evidenced very early in typical

cognitive development. A prediction in my account that seems very plausible given the current evidence is that persons with autism have difficulty comprehending utterances as constative or directive speech acts, whereas children master this communicative dimension from a very early age.

In the last chapter of the book, I will try to show how the naturalistic frame of analysis advocated for constative and directive speech acts can be extended to commissives. I will argue that it is possible to understand why a promise provides the speaker with a desire-independent reason to act, and commits her to do so, without invoking any culture-specific convention or social contract.

1 Austin's distinctions revisited

Before attempting to say what illocutionary acts are, it is important to make sure that we understand what they are not. This is the objective of this chapter. In particular, we will see that it is crucial not to let our theory of illocutions encroach on the territory of *locutionary* acts, nor on that of *perlocutionary* acts. Together with the notion of illocutionary acts, these concepts originate from Austin's (1975) *How to Do Things with Words*. *Perlocutionary acts*, *illocutionary acts*, and – perhaps to a lesser extent – *locutionary acts* are among the most ubiquitous terms of art in pragmatics. Yet, the locutionary/illocutionary/perlocutionary distinction can be easily blurred. Such confusions, as we will see throughout this book, have insidious and unfortunate theoretical consequences.

It is well known that Austin did not have time to revise the notes of his William James lectures, published posthumously as *How to Do Things with Words*, a fact which contributes to make the exegesis of his work notoriously difficult. Fortunately, I need not aim here at a faithful rendering of Austin's views. Austin's observations will constitute a useful starting point for our inquiry. But my ultimate goal is to characterise his three-pronged distinction in terms that help build a psychologically plausible theory of illocutionary force attribution. Therefore, I will often use the insights from contemporary research in order to put up a coherent and plausible analysis, even though the final result might appear to deviate somewhat from Austin.

Actually, we will move some distance from Austin's original views straight from Section 1.1. Austin's distinctions can profitably be viewed as an analysis of communication in terms of the philosophy of action. I will, in a totally anachronistic fashion, examine the levels of meaning marked off by Austin in the light of Davidson's theory of action, which is incompatible with Austin's own take on the topic. Such a strategy will give us a firmer grasp on boundaries between different levels of meaning singled out by Austin. The main points will be as follows:

- (I.a) A phonetic act is the production of a sequence of sounds.
- (I.b) A phatic act is a phonetic act with a morpho-syntactic structure.

- (I.c) A locutionary act is a phatic act with a contextually determined propositional content and a mode of presentation.
- (I.d) Direct and literal illocutionary acts are constituted by locutionary acts, from which they inherit their propositional contents.
- (II) Speaking of perlocutionary acts boils down to describing the causal links between utterances and their effects.

We will see that point (II) provides a neat criterion for disentangling perlocutions from illocutions. As for the notion of locutionary acts, we will not reach a full characterisation before the next chapter. However, towards the end of this chapter I will gather arguments for seeking parallels between locutionary acts and intentional states, and delineate constraints on the definition to be provided in Chapter 2. It is important to spend some time on locutionary acts, because one of the ideas defended throughout the book is that whereas an illocutionary act is *eo ipso* a locutionary act, some locutionary acts have no (direct and literal) illocutionary force.

1.1 Austin's levels of meaning and Davidson's philosophy of action

I will now briefly expose the Davidsonian theory of action I will be using, and explain how I intend to apply it to Austin's locutionary, illocutionary and perlocutionary acts. The exposition of Davidson's views will be sketchy, and I will ignore many important philosophical issues. Even though Davidson's theory of action is widely applied in contemporary linguistics and philosophy it is not beyond criticism (see, e.g. the classic collection of papers by Lepore and McLaughlin, 1985). But my aim here is simply to circumscribe various levels of meaning in order to eventually claim that these distinctions are warranted from a cognitively informed point of view. And, as it happens, Davidson's philosophy of action proves a most handy tool with which to gain theoretical understanding of the relationship between locutions, illocutions and perlocutions. (This being said, matters surrounding action are intricate, and as long as the results of this chapter are left untouched, I am ready to welcome any other framework for making the same points.)

In Austin's opinion (1956, 1966), different descriptions of actions necessarily correspond to distinct ontological entities.¹ Imagine that I deliberately break a teacup; imagine, alternatively, that I break the same teacup by accident. In the first scenario, a description of my action as, say, 'inadvertent' would be false, or, at the very least, infelicitous. The same would apply to the characterisation of the second case as, for instance, 'breaking the teacup on purpose'. This fact prompts Austin to treat the two actions – breaking the cup

¹ Which brings his conception of action very close to Goldman's (1970).

inadvertently and breaking the cup deliberately – as distinct metaphysical types of entities.

However, it is possible that, in both scenarios, the actual physical events are identical: my arm moves across the table, the cup falls down, and breaks on the floor. In a possible world where I break the cup inadvertently and in another one where I break the same cup on purpose, the corresponding physical events could be exactly the same. The core of Davidson's (2001) view is precisely that, along with individuals and objects, our ontology does not contain *actions*, but only *events*, conceived as unrepeatably individuals.² What might look, superficially, like ascriptions of two distinct actions are, actually, alternative descriptions of one and the same event. As a result, event variables³ act as arguments for action predicates. Let us look at this in more detail.

Consider, first, my kissing Nastassia at noon in my office. This action of mine can be truthfully and accurately described in a variety of ways:

- (1) Mikhail kissed Nastassia.
- (2) Mikhail kissed Nastassia at noon.
- (3) Mikhail kissed Nastassia at noon in his office.

But only one single event took place out there in the world: my lips came in contact with Nastassia's. (1), (2) and (3) are just so many ways to describe the same event; they do not denote distinct entities. We thus see that a single event can receive several alternative descriptions. The logical forms corresponding to (1), (2) and (3) would respectively be something like (4), (5) and (6), where *e* stands for one and the same event.⁴

- (4) $Kiss(\text{Mikhail}, \text{Nastassia}, e)$
- (5) $Kiss(\text{Mikhail}, \text{Nastassia}, e) \wedge at\text{-noon}(e)$
- (6) $Kiss(\text{Mikhail}, \text{Nastassia}, e) \wedge at\text{-noon}(e) \wedge in\text{-Mikhail's-office}(e)$

The logical forms in (4), (5) and (6) make clear that (1), (2) and (3) are different descriptions of the same event *e*.

Let us dwell a little longer on cases where several alternative descriptions are available for a single event. An important fact is that, sometimes, whether or not the event at hand can be described in a certain way depends on the

² Here, I am skating over the delicate matter of the individuation of events. Following Quine's (1985) criticism, Davidson abandoned his initial causal criterion (cf. 2001: 163–80) for a definition in spatio-temporal terms (cf. 2001: 309–11).

³ That is, variables that are assigned events as value.

⁴ For ease of exposition, here and below, I implicitly assume that event and individual variables are bound or have been assigned a value.