

1 Introduction

1.1 A conversation at Hondarribia airport

Arriving at Hondarribia airport, John guesses that the pair of students approaching him are in charge of taking him to the pragmatics conference in Donostia. He tells them: “/ninaizdjon/.” Joana, a philosopher, has heard that John is very fond of both identity statements and jokes, so she takes him to have uttered the English sentence:

(1.1) Nina is John.

Although she thinks she has identified the English sentence used, Joana is puzzled about what John could be saying. She expects he is referring to himself with his use of ‘John’.¹ But then to whom is he referring with the typically feminine name ‘Nina’? And why is he saying that he is Nina? What is he trying to do? She suspects John is trying to convey something funny connected with identity sentences and what philosophers say about them, but she can’t figure out what this hypothetical joke might be.

Joana’s friend, Larraitz, a Basque philologist, was not required to learn much about issues of reference and identity, and doesn’t know much about John. This gives her an advantage in understanding what he is saying. She correctly takes John’s utterance to be a use of a Basque sentence,

(1.2) Ni naiz John,

a rather literal if clumsy equivalent of ‘I am John.’ Her only doubts concern the proper way to greet him: shaking hands, the American way, or giving a kiss on each cheek, the usual way in this part of Europe. She decides that she should first respond to John by telling him who she is, using a more appropriate Basque word order than he had:

(1.3) Ni Larraitz naiz [I-Larraitz-am].

¹ We will use single quotation marks for mentioning and as ‘scare’ quotes. We will reserve double quotation marks for utterances, when we don’t number them.

2 Introduction

Both Joana and Larraitz wondered what John had said, and they came to different conclusions. Intuitively, there is a real issue here, with a right and wrong answer; Joana got it wrong, and Larraitz got it right. The issue isn't settled by the sounds that came out of John's mouth. And neither the fact that Basque is the language spoken in the region where the conversation takes place, nor the fact that John's native language is English, settles it either. The answer seems to be provided in large part by what John was *trying* to do, what his intentions were in making the sounds he did. In trying to figure out what John said, a large part of what Joana and Larraitz were trying to do was discover his intentions: what he meant to say. Humans engage in a lot of intentional action, and humans are rather good, given the complexity of the matter, at figuring out why other humans do what they do – at *intention discovery*. Our example suggests that human language and its understanding are an instance of this. Speaking is an intentional activity, and understanding centrally involves intention discovery.

If Larraitz were to try to explicitly reconstruct John's practical reasoning in this case, she would attribute something like the following intentions to him:

- (i) to produce a grammatical sentence of Basque by speaking;
- (ii) to use the sounds appropriate to produce the sentence '[Ni]_{NP} [[naiz]_V [John]_{NP}]_{VP}';
- (iii) to use the indexical 'Ni' to refer to himself;
- (iv) to use the name 'John' to refer to himself;
- (v) to state that *he* is John;
- (vi) to imply that he is ready to trust them to be driven to wherever he is supposed to stay during the conference;
- (vii) to please them by showing that he has learned some Basque.

Intentions do not occur on their own, but with beliefs, and one can have no reasonable hypothesis about intentions without at the same time having a hypothesis about beliefs. In John's case:

- (a) beliefs about the pronunciation and grammar of some Basque expressions, phrases, and sentences;
- (b) beliefs about what these expressions and phrases mean;
- (c) beliefs about his own name;
- (d) beliefs about the correct intonation for assertions in Basque;
- (e) beliefs about what his audience would naturally infer from his assertion and the goal of the conversation;
- (f) beliefs about how Basque speakers are usually very pleased to see a foreigner trying to speak their language.

In order for John to do, by speaking, what he intends to do, his beliefs need to be true or nearly so. If he is wrong about (a) and (b), he may not produce a

meaningful Basque phrase at all. If, for some mysterious reason, he is wrong about his own name, he misleads. If he is wrong about (e) and (f) he won't have the effect on his listeners that he intends, even if he manages to say what he means.

1.2 Three ideas

Pragmatics is the study of how we use language to communicate, and to do the other things we use language to do. Pragmatics involves the formation of intentions on the part of speakers (including, unless noted, writers, typers, and signers) and the discovery of intentions on the part of hearers (including, unless noted, readers and sign interpreters). Pragmatics is but one of the major divisions of the study of language, but it is the one that makes sense of the others. If we look at John's plan, we see it involves knowledge of phonology, syntax, and semantics as well as intentions and beliefs about what he wants to accomplish. That is, John exercises his limited knowledge about what sounds can carry meaning in Basque, how they can be put together into phrases, and what they mean. All of this activity of John's, however, makes sense to us only when we see what he is doing or trying to do; what he wants to say, and to achieve by saying it.

Our approach to pragmatics emphasizes three ideas – we like to think of them as insights – that we think together can provide an approach to the subject that gives a coherent picture of how the parts of language study fit together within a larger picture of human thought and action.

Language as action. The first idea, that we see ourselves getting from Austin, is that language use is a way of doing things with words. Acts of using language, or *utterances*, have a basic structure that is an instance of the general structure of actions: an agent, by moving her body and its parts in various ways, in various circumstances, accomplishes things. By moving my forefinger, in the circumstance in which my hands are poised over a keyboard, I depress the 'j' key;² by doing that, in the circumstance in which the keyboard is suitably attached to a computer and monitor, I make a 'j' appear on the screen. As the example suggests, the circumstances and accomplishments are *nested*; wider and wider circumstances give rise to more and more remote accomplishments. By depressing the 'Return' key, in suitable circumstances, I may send an email that angers a friend, or seals a deal, or precipitates a family or departmental crisis. In our example, John makes noises in circumstances where they count as words of Basque – the circumstances being his intentions and the phonological conventions of Basque – and thereby says something, and thereby, if all goes

² Although the book is co-authored, and there really is no referent for 'I,' we find the first-person singular too effective for presenting examples to give it up.

according to plan, ensures a ride to the conference and pleases the students. As Austin said, by making noises, John performs a locutionary act (saying that he is John), an illocutionary act (introducing himself), and a couple of perlocutionary acts – pleasing Larraitz and puzzling Joana.

Communicative intentions. The second idea, that we see coming from Grice, connects language as action to language as a possessor of content. John said *that he was John*; he conveyed *that he was happy to see the students*. We classify and describe utterances with the same devices, in particular that-clauses of the sort italicized, that we use to describe beliefs, desires, and other mental states. In a wide range of cases, these that-clauses tell us the conditions under which the belief or utterance is true, in others, the conditions under which a desire or request will be satisfied.

Grice's idea was that the meanings of phrases and contents of utterances derive ultimately from human *intentions*, and in particular a special sort of intentions, communicative intentions.³ Communicative intentions have a feature that is characteristic of, but not limited to, the use of language. A communicative intention has its own recognition as one of its goals. John intended to get Larraitz and Joana to believe that he was John, and he intended that an early step in their coming to believe that was to recognize his intention to get them to so believe. From that, together with a bit of common sense and trust, they should conclude that the man speaking to them was indeed John.

Intentions are typically parts of plans. A plan is based on the structure of acts: one does one thing by doing another in certain circumstances. A plan combines a structure of intentions to do one thing *by* doing another with relevant beliefs about the circumstances that support each link. Of course plans can go awry, if the beliefs on which they are based are incorrect. So there are two important structures involved in utterances (and in all intentional action). There is the actual structure: what movements the agent makes, and what results are brought about, given the wider and wider circumstances in which these occur. And there is the structure of the speaker's plan: what movements he intends to make, how he takes the circumstances to be, and so what he intends to accomplish, by moving the way he does, in those circumstances.

Grice's theory of meaning is related to his famous theory of implicatures, which we discuss in Chapter 11, but there is also a certain tension between the two. The tension can be seen by reference to the 'code model' of language. The idea is that the speaker codes up his ideas according to the rules of some language in a linguistic token; the hearer perceives the token, decodes it, and thus comprehends the speaker. Semantics, especially in its more formal versions, can be seen as the key to the code. But the model falls short as an account of

³ He used the term 'M[eaning]-intention' [Grice, 1969a].

1.2 Three ideas

5

all human linguistic communication, for it cannot deal with common linguistic phenomena such as ambiguity, indexicality, and illocutionary force in any straightforward way.

An original motive of Grice's theory of implicatures seemed to be to preserve an important and central place in the theory of language for the coding model and formal semantics. We use the coding model to arrive at what is said; then intention discovery takes over. But Grice's theory of meaning also seems capable of supporting a more radical critique of the code picture. Speaking is a matter of acting on a complex intention; comprehension is discovering the intention; something like coding and decoding may be peripherally involved, but it is not central. We agree with Sperber and Wilson's claim [Sperber and Wilson, 1986] that Grice's picture of meaning and communication can be interpreted either way. We are also sympathetic to their own view that treats intention formation and discovery as central, and decoding as relatively peripheral, at least for understanding ordinary communication in natural language. We return to the issue of coding in the final chapter.

If we combine Austin's and Grice's ideas, we get two results. First, that the *speaker's plan* – what the speaker intends to say and do by making the sounds he does – should be a basic unit of study in pragmatics. It is this that the hearer has to grasp to understand the utterance; it is this that we have to understand, in a way that makes clear how hearers can grasp it, in order to develop our account of pragmatics.

The second result is that what language provides, what all the conventions of Basque and English and the other languages that are recorded in dictionaries and taught to children come to, and what semantics systematically treats, are *ways of acting* and in particular *ways of disclosing one's intentions* to others. Learning language is basically a matter of learning how to do things with words, and in particular how to convey one's own intentions with the help of words, and thereby impart beliefs, desires, suspicions, and all sorts of other things.

Reflexive versus referential truth-conditions. The third idea comes from Perry [Perry, 2000, 2001b]; he thinks of it as ultimately coming from Hume, by way of Perry's work with Jon Barwise on situation semantics [Barwise and Perry, 1983] and with David Israel on information, and Israel and Syn Tutiya on action [Israel and Perry, 1990, 1991; Israel, Perry, and Tutiya, 1993]. This is the distinction between reflexive and referential truth-conditions, and more generally, the idea that utterances and other information-carrying events have different levels of truth-conditions or contents, depending on what one takes as fixed and what one allows to vary. Suppose Elwood falls off his bike, injures his arm, and goes to the emergency room to find out how bad things are. They take an x-ray of his arm, call it **F**. **F** exhibits a certain pattern ψ that shows a break. We might ordinarily say:

6 Introduction

(1.4) **F** shows that the person of whose arm it was taken has a broken arm.

or

(1.5) **F** shows that Elwood has a broken arm.

Hume pointed out that nothing shows anything about the rest of the world, except given some patterns of (more or less) constant conjunction between types of events or situations. He was worried, of course, about how we could legitimately extend such patterns into the future, since our evidence for them holding was based on past observations. We're not worried about that, at least not for the purposes of this book. But we adopt and adapt the idea of (more or less) constant conjunctions as the basis for knowledge gained by observation. Following Barwise and Perry we'll call them 'constraints'; they require that if one type of situation occurs, so does another.

The basic idea is that what an event or a state of a thing – like that pattern on Elwood's x-ray – shows is relative to a constraint. It shows what the rest of the world has to be like, for the event to have occurred, or the thing to be in that state, given the constraint. This conception of the information carried by an event (what is shown by an event; the informational content of an event), makes sense of (1.4), given the constraint:

(1.6) If an x-ray y of a human arm exhibits pattern ψ , then *the person of whose arm y was taken* has a broken arm.

When we instantiate this to **F** and discharge the antecedent, we get (1.4).

This doesn't yet make sense of (1.5), however. Elwood isn't part of the constraint. Most x-rays that exhibit ψ do not show that Elwood has a broken arm, fortunately. Elwood's relevance is that he plays the role, relative to **F**, identified by the antecedent of (1.5). We can say that (1.5) gets at what the rest of the world must be like given first, the constraint (1.6), and, second, the fact,

(1.7) Elwood is the person of whose arm **F** was taken.

Now consider the propositions

(1.8) That the person of whose arm **F** was taken has a broken arm.

(1.9) That Elwood has a broken arm.

Notice that (1.8) is a proposition that is ultimately about **F** *itself*. For this reason Israel and Perry call (1.8) the *reflexive* information carried by **F**, or, more correctly, carried by the event or fact that **F** exhibits pattern ψ .⁴ In the terminology of this book, we could call it the *x-ray-bound* information. Suppose

⁴ They also sometimes call it the 'pure' information, for reasons lost in the fog of history.

1.2 Three ideas

7

a nurse found the x-ray on the floor, with no indication of whose arm had been x-rayed. She would know that the person it was of had a broken arm, but she would have no way of identifying that person except as the person of whom it was taken; her knowledge is *bound*, in that sense, to the x-ray **F**.

(1.9) on the other hand is not about **F** but about Elwood. It could be true even if **F** were never taken; Israel and Perry call (1.9) the *incremental* information carried by **F**, by which they meant that it gets at what the rest of the world must be like, for the x-ray to turn out as it did, given not only the constraint (1.6) but also, in addition, the ‘connecting fact’ (1.7).

In *Reference and Reflexivity* and elsewhere, Perry argues that basically the same distinction needs to be made, in the philosophy of language, with respect to the truth-conditions or content of utterances. The idea is that we can consider under what conditions an utterance might be true – or more precisely, what the rest of the world has to be like for the utterance to occur and be true – simply given the constraints on truth-conditions provided by the meanings of words as fixed by the conventions of language, or, taking in addition, further facts about the utterance, such as the speaker, time, place, and objects referred to with the use of names and demonstratives. Suppose Elwood says

(1.10) I have a broken arm.

Call his utterance **u**.

The rules of English tell us that

(1.11) Any English utterance **u** of the form ‘I have a broken arm’ is true if and only if the speaker of **u** has a broken arm.

If we instantiate **u** and discharge the antecedent, we get, as the truth-conditions of **u**:

(1.12) That the speaker of **u** has a broken arm.

That’s what Perry calls the ‘reflexive truth-conditions’ of **u**, as we do in this book, but we also call it the ‘utterance-bound’ truth-conditions. If, in addition, we are given that fact,

(1.13) Elwood is the speaker of **u**

then we get (1.9) as the truth-conditions of **u**. Perry calls this the ‘referential’ truth-conditions, or content, of **u**, as do we in this book.

Given the complexity of language, and especially the different sorts of roles that the rules of language establish, the simple distinction between utterance-bound and referential truth-conditions gives way to a more complex scheme, as we shall see in the ensuing chapters.

1.3 The anatomy of an utterance

Now let's return to John's utterance. John's act is in a sense rather simple; he simply produces the sound '/ninaizdjon/.' And yet by doing this he manages to do something rather complex; there is a lot of structure in his plan of action. He intends to make certain sounds, and thereby produce certain phonemes of Basque, and thereby say something, and thereby have various effects on his listeners. The sound '/ninaizdjon/' makes only an incremental contribution to accomplishing all of this; John relies on a lot of structure that is already in place. He relies on the conventions of Basque, on the structure of the particular situation, in which the students will see him as using language to say something to them, and he relies on the structure of human psychology and Basque culture. John's utterance is hardly a very dramatic example; wars have been started, discoveries promulgated, lives saved, and philosophical positions refuted by pronouncing a few suitable syllables in the right situations.

John realizes that the effects he wants to have on Joana and Larraitx will not be produced merely by their hearing the sounds '/ninaizdjon/.' He wants them to know that he is John; they will do this by recognizing that he is saying he is John (and trusting that he is sincere and knows who he is). They will recognize what he is saying by recognizing what he is trying to say. If John's pronunciation is too terrible, or his syntax too garbled, he might try to say that he was John and not manage to say it. And yet Larraitx might still have figured out what he was trying to say, and that's what would have been important in her coming to believe that he was John. But of course, usually the best way to convey to people what you are trying to say is to successfully say it.

John also wants to convey to Joana and Larraitx that he is happy to put the next leg of his journey in their hands. This fact does not follow from the fact that he is John. Indeed, he might have made his pleasure and relief clear without using language at all, simply by looking pleased and entrusting his luggage to them to carry to the car. His communication, though not involving language, still fits the Gricean model; his intention is not simply to get help with his bag, but to let them know he is pleased to be in their hands, by their recognizing that he intends to convey this.

Finally, John wanted to please Joana and Larraitx by his efforts to learn a little Basque. Just saying that he is John won't get this effect. He has to say it in a certain way – in Basque, not in English, and in good enough Basque to suggest he has made some effort. Here the intention is not Gricean. John wants to please them, but not necessarily by recognition of his intention to do so. He wants to please them because they are impressed with the quality of his Basque and the effort that must have gone into learning even that much of this notoriously difficult language. This is what Austin calls a perlocutionary effect,

1.3 The anatomy of an utterance

9

a result of one's saying what one does, but a further effect, not something one does *in* saying it.

Joana's and Larraitx's task is then, in a sense, quite formidable. Their challenge is to infer, from the few sounds they hear, the complex of intentions, the plan of action, that animates John. Just as John could not hope to accomplish what he plans without relying on pre-existing structures, they have no hope of discovering his intentions without building on a lot of structure: the conversational situation, the larger situation involving the conference, human psychology, John's psychology, and the conventions of language. Larraitx meets the challenge, while Joana does not.

Joana's problems began with a wrong hypothesis about John's overall purpose in speaking – what he was trying to do by saying what he did. This is what we call *far-side pragmatics*, that is, pragmatic reasoning about why the speaker says what he does, what he is trying to communicate or accomplish by saying it. This mistake of Joana's led to a mistake about *near-side pragmatics*, that is, pragmatic reasoning about what the speaker is trying to say by producing the sounds he does. If she hadn't suspected that John was trying to be funny, then when she saw that the apparent English sentence, (1.1), made no sense, she might have questioned her assumption that he was speaking English, and recognized his less than fluent Basque for what it was.

The near-side and far-side terminology suggests that there is a central aspect of an utterance that marks an important divide. Historically, this aspect has been identified with 'saying something,' and we'll stick with that terminology for now, although in time we'll use Austin's concept of a *locutionary act* to explicate it. The picture is that a speaker produces sounds *in order to say something*, and says something *in order to accomplish further results* – to communicate further information, or to perform various speech acts. The sorts of knowledge and planning that the speaker has to bring to bear on the near side to get something said have largely to do with the conventions of language. He needs to know that certain sounds are ways of producing certain phonemes in a given language; that producing a certain string of phonemes is a way of producing a certain sentence in that language; and that in that language and in the context he is in, producing that sentence is a way of saying what he has chosen to say. His reasoning on the far side, about what he will accomplish in or by saying what he does, has mostly to do with people, context, and culture rather than the conventions of the language he uses. John drew on his meager knowledge of the conventions of Basque to say that he was John; he drew on his slightly more extensive knowledge of Basques to realize that by saying this, he would communicate his pleasure in having the students take charge of his journey, and please them with his efforts to speak their language.

1.3.1 Far-side traditions

Both Austin and Grice focus on what speakers try to do beyond saying what they do; Grice on their further communicative intentions, Austin on the speech acts they intend to perform in or by saying it. From the 1960s until the mid 1980s, with some important exceptions, such issues of far-side pragmatics dominated the field of pragmatics.

Grice was concerned with making a distinction among the contents conveyed by an utterance – within the ‘utterance meaning,’ in his terminology. He distinguished between what a speaker *says* and what she *implicates* in uttering a sentence. He convincingly showed that there always are contents that a speaker communicates without saying. The speaker does not codify them in sentences, so they are not there for the hearer to decode. The hearer has to infer them, attending to pragmatic principles and contextual information aimed at recognizing the speaker’s intentions, because the speaker’s communicative intentions are intended to be recognized by the addressee. Grice thus showed that there is much more to human communication than using semantics of language to code up one’s thoughts in language, and decode the utterances of others. Pragmatic reasoning is also critical for a theory of human linguistic communication. This was particularly clear for the case of implicatures. In our example, intentions (i–v) and beliefs (a–d) would be the relevant ones for determining what John said; the remaining ones would affect what he implicated by his saying it and other perlocutionary aspects of his utterance.

Austin was concerned with the difference between what a person said and other ‘speech acts’ that he performs, and, in this latter class, between those things he does *in* saying what he does (the ‘illocutionary acts’) and those things he does *by* saying what he does (the ‘perlocutionary acts’). So he made a threefold distinction among the different levels of a speech act, so that intentions (i–iv) and beliefs (a–c) would correspond to the *locutionary* act performed, all those plus intention (v) and belief (d) would determine the illocutionary act performed, and the remaining one would be relevant to the (intended) *perlocutionary effects* of the speech act. Austin’s threefold distinction and Grice’s twofold distinction are not competitors, but complement one another. In their broad lines, at least, most philosophers and linguists accept both sets of distinctions.

Hence both Grice and Austin saw a major theoretical break between what is said (or the locutionary act) and the further things accomplished in speaking. Study of that latter is the sort of pragmatics we call ‘far-side pragmatics.’ If this line is identified with the border between semantics and pragmatics, far-side pragmatics is all there is to the discipline.

1.3.2 Near-side debates

But as our example makes clear, far-side pragmatics does not exhaust pragmatics. That is, intention and intention-discovery are involved on the near side,