

Joint Meaning

Antonella Carassa¹, Marco Colombetti¹²

¹ Faculty of Communication Sciences, University of Lugano, Lugano, Switzerland

² Department of Electronics and Information, Politecnico di Milano, Milano, Italy

Abstract

In this paper we want to reconcile two apparently conflicting intuitions: the first is that what a speaker means is just a function of his or her communicative intentions, independently of what the hearer understands, and even of the actual existence of a hearer; the second is that when communication is carried out successfully, the resulting meaning is, in some important sense, jointly construed by the speaker and the hearer. Our strategy is to distinguish between *speaker's meaning*, understood as a personal communicative intention, and *joint meaning*, understood as a joint construal of the speaker and the hearer. We define joint meaning as a type of propositional joint commitment, more precisely as the joint commitment of a speaker and a hearer to the extent that a specific communicative act has been performed by the speaker. Joint meaning is therefore regarded as a deontic concept, which entails obligations, rights, and entitlements, and cannot be reduced to epistemic and volitional mental states like personal belief, common belief, personal intention, and communicative intention.

1. Introduction

In the last fifty years the theory of linguistic meaning has witnessed an impressive sequence of major developments. Until the mid Twentieth century, meaning has been regarded as a property of sentences or, more generally, of signs. Since then, this view has been completely revolutionized by the contributions of several authors, like Grice (1957), Austin (1962), Strawson (1964), Searle (1969, 1979), Schiffer (1972), Bach and Harnish (1979), and many others.

One of the main achievements of this research tradition, which we shall simply refer to as *Speech Act Theory*, has been the development of a conceptual framework in which theories of meaning and communication can be stated. In this framework the use of language is analyzed in terms of speech acts, performed by subjects entertaining certain epistemic and volitional mental states, like personal beliefs, common beliefs (also known as common ground), personal intentions, and communicative intentions. Meaning is regarded as a function of the speaker's mental states, and in particular of his or her communicative intentions.

More recently, the approach of Speech Act Theory has been criticized by Herbert Clark, who considers it incapable of accounting for the intrinsically social nature of communication. In *Using Language* Clark (1996:130) submits a view of meaning as a joint construal of the speaker and the hearer:

“The principle I wish to defend is this:

Signal recognition principle. Signaling and recognizing in communicative acts are participatory acts.

The joint act of one person signaling another and the second recognizing what the first meant I will call a *communicative act*.”

Clark (1996:131) goes on defining the meaning of a communicative act r as follows:

“Speaker’s meaning (joint)

In presenting [signal] s to A , speaker S means for A that p if and only if:

0. the communicative act r includes 1 and 2;
1. S presents s to A intending that p as part of r ;
2. A recognizes that p as part of r .”

In this definition it is unclear whether Clark assumes speaker’s meaning to include the actual recognition by A that p is part of r (as it seems from a literal interpretation of the definition), or rather the speaker’s intention that A recognize that p as part of r (as it appears from other parts of Clark’s book). In any case, what Clark wants to do is to build on the intuition that communication involves some sort of cooperation of the speaker and the hearer.

In this paper we want to reconcile two apparently conflicting intuitions: the first is that what a speaker means is just a function of his or her intentions, independently of what the hearer understands, and even of the actual existence of a hearer; the second is that when communication is carried out successfully, the resulting meaning is, in some important sense, jointly construed by the speaker and the hearer. Our strategy will be to distinguish between *speaker’s meaning*, understood as a personal communicative intention, and *joint meaning*, regarded as a joint construal of the speaker and the hearer.

To develop our proposal, we shall partly build on some of Clark’s own ideas, like his concept of a joint project and the related notions of joint purpose and joint commitment. However, we contend that an adequate treatment of such concepts requires broadening the classical view of mental states as consisting of epistemic or volitional attitudes. More precisely, we shall argue that human interactions have an irreducible deontic component, which should be carefully considered by theories of meaning and communication. To model such a component we shall rely on the concept of commitment, and in particular of joint commitment as developed by Gilbert in several articles and books (1989, 1996, 2000, 2006).

The importance of deontic concepts in general, and of commitment in particular, for modeling human communication has been recognized long ago.¹ In their pioneering book, *Understanding computers and cognition*, Winograd and Flores (1986:76) wanted “to counteract the forgetfulness of commitment that pervades much of the discussion (both theoretical and commonplace) about language”; in dialogue theories, commitment-based models have been proposed and discussed since the concept of a commitment store was introduced by Hamblin (1970) and later developed by Walton and Krabbe (1995); and, more recently, Searle (2007) has advocated a view of human

¹ Interestingly, a deontic perspective is already embedded in the etymology of the word *communication*. “To communicate” means making something common; in turn, “common” comes from the Latin *communis*, composed by *cum* (with) and *munus*, which means both *gift* and *duty* (especially a duty associated to an institutional position). If duty is obviously a deontic category, so is gift: in the Latin world, as in many cultures, gifts were part of a complex system of social practices subject to strict obligations.

language in which deontic normativity is regarded as a constitutive component of human language, side by side with representative power and syntactic compositionality. The deontic side of interaction has also been a major topic of interest in the field of multiagent systems (Singh, 1998; Colombetti, 2000; Grosz and Hunsberger, 2006), and is currently becoming an important issue in theories of human intersubjectivity (Carassa et al., 2008). In our opinion, however, none of the proposals put forward so far gives a satisfactory account of the deontic component of meaning and communication. In this article we aim at providing such an account; in particular we shall give a deontic definition of joint meaning, which we submit as an adequate basis for an analysis of communicative interactions.

The paper is structured as follows. In Section 2 we clarify our theoretical background, specifying the elements of classical Speech Act Theory that we take for granted, introducing Clark's concept of a joint project, sketching Gilbert's theory of joint commitments, and integrating it with additional elements from different sources. In Section 3 we present our view of speaker's meaning, and analyze the relationships between communicative intentions and commitments. In Section 4 we introduce and discuss our concept of joint meaning. Finally, in Section 5 we draw some conclusions and suggest directions for future research.

2. Theoretical Background

To clarify our main theoretical assumptions, we shall review the elements of Speech Act Theory that we take for granted, Clark's view of communicative acts as participatory actions, and Gilbert's theory of joint commitment.

2.1 Speech Act Theory

Speech Act Theory is by no means a monolithic field. In particular, the two best established accounts—those provided by Searle (1969) and by Bach and Harnish (1979)—differ substantially in the very conception of what a speech act is. However, at least on two points all accounts do agree. The first is that communicating is a matter of performing certain types of actions, called *speech acts*, which typically (but not necessarily) involve the use of language. The second is that three different kinds of speech acts can be distinguished. At a lower level we find certain actions (Austin's *locutionary acts* or Searle's *utterance acts*) that are instrumental to the realization of higher level actions, called *illocutionary acts*; in turn, an illocutionary act is typically a means for performing a *perlocutionary act*. There is also general agreement on the fact that illocutionary acts are the carriers of fully fledged meaning.

As far as illocutionary acts are concerned, two competing views can be identified. According to the first view, already submitted by Austin (1962), performing an illocutionary act is a matter of exploiting certain conventions. Searle's theory of illocutionary acts can be regarded as a development of this view, according to which illocutionary acts are made possible by the existence of suitable constitutive rules, which are part of the institution of language. According to the second approach, first proposed by Strawson (1964) and then developed into a comprehensive theory by

Bach and Harnish (1979), illocutionary acts are to be classified into two categories: while the acts of one category, usually called *conventional illocutionary acts*, do require specific conventions, the acts of the other category, called *communicative illocutionary acts* or simply *communicative acts*, do not. Producing and understanding a communicative act is then a matter of the speaker having, and the hearer recognizing, certain configurations of intentions, often called *communicative intentions* or *Gricean intentions*, as they have first been described by Grice (1957) in his famous paper on nonnatural meaning; a communicative act is therefore successfully performed if *uptake* is secured, that is, if the corresponding communicative intention is recognized. As these concepts are going to play an important role in our definition of joint meaning, we shall now expand on them.

Communicative intentions

As is well known, Grice (1957:385) defined nonnatural meaning, or meaning_{NN}, as follows:

- (1) “«A meant_{NN} something by x » is (roughly) equivalent to «A intended the utterance of x to produce some effect in an audience by means of the recognition of this intention»”

A problem of Definition 1 is that it is not clear what the expression “this intention” refers to. A possibility is that the expression refers to the speaker’s intention “to produce some effect in an audience” through his or her utterance. This interpretation involves two levels of intention: the first-level intention to produce an effect in an audience, and the second-level intention that the first-level intention be recognized by the audience. Another possible reading of the definition is that “this intention” refers to the whole intention “to produce some effect in an audience by means of the recognition of this intention.” This interpretation involves a single reflexive intention, which can be described as: (A intended the utterance of x to produce some effect in an audience by means of the recognition of *this*), where “*this*” refers to the whole expression in parentheses. The two-level interpretation has been taken, if only as a starting point for further analysis, by Strawson (1964) and by Searle (1969), and seems to be coherent with Grice’s later writings (e.g., Grice, 1969); the reflexive interpretation has been preferred, for example, by Bach and Harnish (1979).

Historically, there is little doubt that Grice had in mind a two-level, nonreflexive interpretation; this is proved by the fact that he did not object to the nonreflexive readings of his definition provided by Strawson and by Searle (see also Kuroda, 1989, for a defense of the two-level interpretation). However, it can be argued that a concept of communicative intention is better captured by a reflexive definition. To see that this is the case, first consider that Strawson (1964) already showed that two levels of intention are not sufficient to rule out certain counterexamples, which intuitively do not involve communication but satisfy a two-level interpretation of Grice’s definition. Strawson proposed a three-level definition, in which also the second-level intention is intended to be recognized. However, similar (although more complex) counterexamples can be shown to undermine Strawson’s proposal, as well as any definition based on a finite number of levels of intentions (see for

example Airenti et al., 1993); only a reflexive interpretation of communicative intentions seems to match our intuitions about communication.

Grice's definition has also been criticized, in particular by Searle (1969), because it has been viewed as describing a perlocutionary, rather than illocutionary, act. Contrary to this interpretation, it can be argued that Definition 1 contains *in nuce* all basic elements of a theory of speech acts. Indeed, Grice's definition regards the production of meaning as a type of action, because it describes it as the intentional attempt to bring about a result. Moreover, the definition introduces a clear distinction between: an instrumental act, x , which one can see as a locutionary or utterance act; an intended effect on an audience, which can be viewed as a perlocutionary effect; and the intended recognition of the speaker's intention, which can be regarded as an illocutionary effect, as far as we limit our attention to communicative illocutionary acts.

Another important point concerns how we should conceive of the hearer's recognition of a communicative intention. As argued by Schiffer (1972), recognition cannot be simply equated with a personal belief of the hearer: for a communicative act to be successful, it is crucial that the speaker's communicative intention becomes common belief of the speaker and the hearer.

We think that these considerations justify taking Definition 1 (in its reflexive reading, and interpreting recognition as common belief) as a starting point for an analysis of communication. However, two important issues have to be addressed: the first is to clarify what it means to achieve an effect in an audience "by means of" something; the second is to explain how the recognition of a reflexive intention can actually work as a means to achieve an effect in an audience. We shall come back to these points in Section 3.

Uptake

In *How to do things with words*, Austin (1962:117) remarks that "the performance of an illocutionary act involves the securing of *uptake*." Probably, Austin wanted to say that an illocutionary act, conceived as a conventional act, may be successful only if it is recognized by the hearer as being performed according to the relevant conventions. Strawson (1964), however, argued that uptake is crucial only in the case of communicative illocutionary acts, and defined uptake as the recognition of the speaker's communicative intention; this view then became part and parcel of the theory of communicative acts developed by Bach and Harnish (1979).

In *Using language*, Clark is very critical of the 'individualism' of classical Speech Act Theory. In particular, he remarks that according to Searle's conception of a speech act (Clark, 1996:137),

"It doesn't matter whether anybody receives, reads, or understands it. This view is, of course, absurd. [...] Austin recognized the problem, but his suggestions were ignored by most who followed."

Clark's idea of uptake considerably departs from Strawson's, as he conceives of uptake as going beyond the mere understanding of a speaker's communicative intention. In Clark's view communication is a means to propose, negotiate, and carry

out *joint projects*, which can be *taken up* (i.e., accepted) or not by a hearer: uptake is not limited to understanding, but involves accepting a joint project.

Clark's conception of uptake may be inappropriate, because it conflates aspects that are part of illocution with aspects, like the actual acceptance of a joint project, that pertain to perlocution. However, Clark is probably right in suggesting that uptake involves more than the recognition of a communicative intention: intuitively, there is a sense of participation in the concept of uptake of which the received view of Speech Act Theory does not account. Between the classical and Clark's conceptions of uptake, however, we believe that there is a third, interesting possibility: regarding uptake as a kind of participation that still does not go as far as accepting a joint project. This is the view that we shall develop in Section 4.

2.2 Clark's view of communicative acts

According to Clark, the meaning of the communicative act produced by a speaker (that we shall call the *initiator* of the exchange) appears to be collectively construed by the initiator and by his or her partner. Suppose for example that Barry (the boss) has the following exchange with Evelyn (an employee):

- (2) Barry: Are you free tonight?
 Evelyn: Yes, I am, Barry. I think it would be a good idea to spend a couple of hours on the draft of the project.
 Barry: Oh, well ... yeah. Back here at half past eight?
 Evelyn: Perfect, I'll be here.
 Barry: See you at eight thirty, then.

Now imagine that Barry's original goal was to invite Evelyn to dinner at a nice restaurant, and that his original question was just a pre-request (Clark, 1996) in this direction. In such a case we can say that Evelyn redefined Barry's original intention as an indirect proposal to spend the evening working, and that such a redefinition has been accepted by Barry. Using Clark's terminology, the proposal to spend the evening working on the draft must be regarded as a *joint construal* by Barry and Evelyn.

Joint construal is pervasive in communication, in particular when a communicative act leaves room for reinterpretation, like in the case of indirect illocutionary acts. Pretheoretically, it is tempting to interpret instances of joint construal as cases in which the meaning of a communicative act is 'negotiated' by the initiator and his or her partner. However, this brings in a problem for Speech Act Theory. On the one hand, if speaker's meaning is defined in terms of the speaker's intentions, it would be inappropriate to assume that such meaning can be negotiated: although intentions are ontologically subjective, in the sense that they literally are states of a subject, they are epistemically objective, in the sense that they are objective facts of the matter (Searle, 1995); this implies, for example, that Barry's actual communicative intention when he uttered "Are you free tonight?" is not subject to external arbitration. On the other hand, in Exchange 2 something akin to negotiation clearly goes on. So the problem is, *what* is negotiated, if it is not speaker's meaning?

It is interesting to note that Exchange 2 is open to different interpretations. For example, one may assume that Evelyn misinterpreted Barry's pre-request, and is then sincere in showing that she understood it as an indirect proposal to spend the evening working together. A second option is that Evelyn correctly understood that her boss was about inviting her for dinner, but decided to change his tune. In the latter case, the boss may have understood Ann's maneuver to circumvent the invitation, or may believe that she really misunderstood his original question. But while only Barry and Evelyn know what went on privately in their minds, the joint construal of Exchange 2 as an accepted proposal to spend the evening working is public, not private, and any theory of meaning must account for this.

Clark's treatment of joint construal must be considered in the context of his larger theory of linguistic communication. According to Clark (1996:191), "Signaling is of interest only because it is used in advancing the joint activities people are engaged in"; communication is therefore viewed as a means to propose, negotiate, and carry out a joint project. A clear distinction is drawn between *construing* a joint project and *committing* to it; more precisely, Clark argues that joint projects serve joint purposes, which in turn require the commitment of all the participants. Commitment to a joint purpose is defined as follows (Clark, 1996:203):

- (3) "For A and B to commit to joint purpose *r*:
1. *Identification* A and B must identify *r*
 2. *Ability* It must be possible for A and B to do their parts in fulfilling *r*
 3. *Willingness* A and B must be willing to do their parts in fulfilling *r*
 4. *Mutual belief* A and B must each believe that 1, 2, 3, and 4 are part of their common ground"

In our opinion, this definition has a substantial problem, because joint commitment cannot be defined solely in terms of epistemic and volitional mental states like personal beliefs, mutual beliefs, and intentions. To see that this is the case, let us elaborate on a possible development of Exchange 2. By carrying out the exchange Barry and Evelyn not only define a joint project, but also commit to its joint purpose. Now suppose that in the interval of time between the end of the exchange and eight thirty Evelyn changes her mind, and is no longer willing to do her part in fulfilling the joint purpose. According to Definition 3, the joint purpose just ceases to exist, and there is nothing that Evelyn ought to do to negotiate a new course of interaction. In other words, given that a joint commitment depends only on epistemic and volitional mental states, it can be unilaterally rescinded without any further consequence. Clause 4 of Definition 3, which brings mutual belief into play, does not change the situation: when Evelyn changes her mind, Barry will just go on believing that every relevant component of the joint commitment is common ground, while Evelyn will stop believing so; but this, by itself, need not have any consequence on Evelyn's behavior.

In real interactions, however, things go on differently. After changing her mind, Evelyn is likely to contact Barry to cancel the appointment; she is also supposed to provide some justification, like a sudden headache or the unforeseen necessity to

visit her mother. In any case, there is something that she *ought to do*. The problem with purely epistemic/volitional accounts of joint commitment, like the one put forward by Definition 3, is that they do not imply any *ought to* clause.

In a recent work, Clark (2006) extends his analysis of communicative interactions in terms of joint commitments and argues, quite rightly in our opinion, that two crucial features of joint commitments are that they constrain an agent's autonomy and that they persist through time. However, Clark offers no analysis of the logical structure of joint commitment. Indeed, a detailed analysis of this concept, together with the related notion of a *plural subject*, has been proposed and extensively analyzed by Margaret Gilbert in several books and papers (see in particular 1996, Part III; 2000, Chapter 4; and 2006, Chapter 7). In the next subsection we briefly describe the main features of Gilbert's proposal.

2.3 Gilbert's theory of joint commitment

According to Gilbert all genuinely collective phenomena (like joint activities, collective beliefs, group feelings, social conventions, and so on) involve a special kind of commitment, that she calls a *joint commitment*. A subject may be personally committed to do X, for example as a result of an individual decision: such a decision may be rescinded, but until this does not happen the subject is committed to do X. Being committed to do X is a reason (although in general not a sufficient cause) for the subject to do X; however, in the personal case the subject is the only 'owner' of the commitment, and can rescind it as he or she pleases. Contrary to personal commitments, a joint commitment is a commitment of two or more subjects, called the *parties* of the joint commitment, to engage in a common project "as a single body." Taken together, a number of subjects jointly committed to X-ing as a body (where X may be doing something, or believing something, or intending to do something, and so on) form a *plural subject* of X-ing. The main difference between personal and joint commitments is that joint commitments are not separately owned by their parties, but they are, so to speak, collectively owned by all parties at the same time.² Joint commitments may arise as a result of explicit agreements, but this is not strictly necessary. According to Gilbert, to create a joint commitment it is necessary and sufficient that every party expresses his or her readiness to be so committed, in conditions of common knowledge; such common knowledge may derive from an explicit agreement (like in the case of Exchange 2), but also from less structured interactions and, in many cases, from the shared understanding of a culturally meaningful context.

² Gilbert's definition of plural subjects has been alleged of being circular (e.g., by Tollefsen, 2002), because the notion of X-ing "as a body," used to define plural subjects, seems to presuppose the very concept of a plural subject. In our opinion, however, the definition involves no circularity: what Gilbert appears to assume is just that human beings have a primitive mental capacity to understand a group of people as doing something together. Similar assumptions have been adopted in other accounts of collective intentionality, like for example Searle's (1990).

In view of the purpose of this article, the main feature of joint commitments is that they consist of deontic relationships between the parties, like directed obligations and the correlative rights and entitlements. If a group is jointly committed to do something, then every party is obligated to all other parties to do his or her part of the joint activity, and has the right that all other parties do their parts. It is characteristic of joint commitments that such obligations are *created simultaneously*, and are *interdependent* in the sense that each party is bound by the joint commitment only as long as the other parties are so bound. If and when all its obligations are fulfilled, a joint commitment is itself fulfilled; on the contrary, if one of its obligations is violated, in many cases a joint commitment ceases to exist, and all parties are freed of their obligations.

2.4 Social commitments

Gilbert's theory of joint commitment does not cover all types of commitment that are relevant for understanding communication. Indeed, there seem to be different types of commitments. Clark (2006), for example, mentions four different types: private self-commitment, when a subject is committed to him or herself to do an action, without letting anyone know; public self-commitment, when one makes a commitment to him or herself in front of other people; simple other-commitment, for example when one makes a unilateral promise to another person; and participatory commitment, when a group of people jointly commit to carry out a common project. In what follows we shall neglect self-commitments, which are not relevant for our current goals, and shall concentrate only on commitments that involve more than one subject, which we call *social commitments*, or simply *commitments* for conciseness' sake.

Social commitments

Social commitments are desire-independent reasons for action that, contrary to other types of desire-independent reasons for action like moral or legal obligations, are intentionally created by subjects. They can be viewed as bundles of deontic relationships, intentionally created by the very subjects that are connected by such relationships. A unilateral promise, for example, creates a directed obligation of a *debtor* to a *creditor*, because the two parties so agreed.

Joint commitments are an important subclass of social commitments. They arise in case of agreements, but also (according to Gilbert, 2006) in the case of unilateral promises. Suppose for example that Albert promises to Barbara to climb the Mount Everest. While carrying out the promised deed will be a personal action of Albert, and not a joint project of Albert and Barbara, it seems reasonable to regard Albert and Barbara as jointly committed to the fact that they are now bound by the promise. This involves also directed obligations of Barbara to Albert, like for example the obligation not to hamper Albert's efforts to climb the Mount Everest. Moreover, as Gilbert suggests, none of the parties has the power to unilaterally rescind the promise, and this implies that Albert's action is now 'owned' by both parties.

There are, however, types of social commitments that do not appear to be cases of joint commitment. Suppose for example that Albert tells Barbara that he has already climbed Mount K2. Albert is now committed to Barbara to the truth of his assertion (in a sense that will be clarified below, when we discuss propositional commitments); but it seems to us that there is no sense in which Albert and Barbara are *jointly* committed to the fact that Albert climbed Mount K2.

The function of social commitments

By making commitments, people intentionally constrain their own wills. But why should they do so? Why should they put limits to their own freedom to act? In previous works (Carassa et al., 2008) we argued that, by doing so, people bring a sort of stability to interactions that would be impossible to achieve otherwise.

It is generally accepted that humans are the only species that can deploy a very specific type of rationality, that is, the ability to plan their future. As anticipatory planning is one of the distinctive features of *Homo sapiens* (see for example Gärdenfors, 2008), it is not surprising that so much attention has been devoted to it in Cognitive Psychology, Philosophy of Mind, Economics, and Artificial Intelligence. The function of future-directed intentions, or prior intentions in Searle's terminology (1983), has been analysed, among others, by Michael Bratman (1987), who stresses their characteristic role of coordinating practical reasoning. Future-directed intentions, organized into complex plans, allow human subjects to reason within stable tracks directed to specific purposes, thus avoiding the risk of being misled by fluctuating motivations.

Social commitments, and joint commitments in particular, play a stabilizing role with respect to human interactions that is analogous to the role played by future-directed intentions in the case of individual actions. They achieve this function by creating directed obligations, thus decoupling future actions from possibly fluctuating motivations. Consider the following example: by making a suitable agreement, Sarah and David jointly commit to paint their apartment together. While the joint commitment is in force, Sarah and David are obligated to each other to carry out appropriate actions, like helping one another to carry out their collective activity in face of possible difficulties. Given such a commitment, that Sarah and David maintain their personal motivation to paint the apartment is not crucial, because the reason for doing so is now an obligation of their joint commitment.

Being reasons for action, social commitments have an essential relationship with intentions, which can be described as follows. As we have seen, a commitment to do action α logically entails certain obligations, and in particular the debtor's obligation to do α ; in turn, such obligation is a practical reason to form the intention to do α ; for example, the promise to quit smoking entails the debtor's obligation to quit smoking, and such an obligation is a practical reason for him or her to quit smoking intentionally. However, a commitment to do α does *not* logically entail the debtor's intention to do α ; in other words, it is logically possible for a subject to be bound by a commitment to carry out a course of action and still form no intention to do so. In such a case, of course, the commitment will eventually be violated.

Propositional commitments

The content of a social commitment need not be an action type: a subject may commit to a state of affairs, and in particular to a mental state, like a belief or an intention. The dual nature of commitment, as either an *action commitment* or *propositional commitment*, has been analyzed by Walton and Krabbe (1995) in the context of their theory of dialogue. These authors take action commitment as the only primitive concept, and define propositional commitment as derived from action commitment. More specifically, Walton and Krabbe define the commitment to a proposition as the commitment to perform certain actions related to such proposition, depending upon the context. To give an example, suppose that in a discussion Albert commits to the fact that he is a devote Roman Catholic. If later on Albert enters a Catholic church, he will have to behave in certain ways, or else his propositional commitment will be violated.

It is important to remark that all types of commitments, inclusive of propositional commitments, are desire-independent reasons for *action*. Providing a general characterization of what courses of action would violate a propositional commitment is a difficult task, which we are not going to tackle in this paper. We shall simply assume that in every specific context it is possible to establish whether a personal or collective action is 'required' by a proposition, in the admittedly intuitive sense that failing to carry out that action would, in the relevant context, be incompatible with the proposition. In this way, a commitment to a proposition is regarded as a mechanism for the generation of action commitments, whose contents are given by the context-dependent actions that are required by the proposition.

Often, propositional commitments are commitments to having an attitude of a given type, including beliefs, desires, intentions, dispositions, feelings, and so on. For example, in saying that he likes fish, George commits to having a certain type of attitude relative to eating fish. It is important to remark that fulfilling or violating a commitment to an attitude is logically independent of sincerity. If a subject acts coherently with a commitment, then the commitment is fulfilled even if the subject was insincere about his or her beliefs, desires, and so on; and if a subject does not act coherently with a commitment, then the commitment is violated even if the subject was sincere. Of course there are often moral reasons for being sincere, but this has nothing to do with the obligations of commitment.

3. Speaker's Meaning

As we said in the Introduction, our goal is to reconcile two apparently conflicting approaches to meaning. On the one hand, the Gricean tradition regards meaning as a function of the speaker's intentions, independently of what the hearer actually understands; on the other hand, there is an undeniable appeal to Clark's view that meaning is, in some important sense, jointly construed by the speaker and the hearer. We propose to solve this apparent dilemma by distinguishing between *speaker's meaning*, understood as a function of personal communicative intentions, and *joint meaning*, understood as a joint construal of the speaker and the hearer. Here we shall

develop our approach to speaker's meaning, and shall come to joint meaning in the next section.

From Speech Act Theory we accept that communication is a matter of achieving an effect in a hearer through the performance of a communicative act, whose successful performance involves securing uptake, which in turn requires the recognition of reflexive communicative intentions. But how exactly do we achieve effects on people? And why is the recognition of a reflexive intention a necessary component of communication?

3.1 Bringing about effects in people

If we regard a person as a mere physical object, we can produce effects by acting physically on his or her body, for example by applying forces, by injecting chemicals, by heating or freezing, and so on. One may also exploit low-level physical interactions to produce mental effects; for example, one may hurt a person to cause pain, make noise to cause auditory perceptions, or bring an object in a subject's visual field to cause a visual perception. Clearly, this is not what is at stake when we deal with communication: a communicative effect on a person's mental states is not produced by sheer physical causation.

It seems to us that there are two distinct ways of producing a communicative effect on mental states: by providing *reasons*, and by providing *affordances*.

Providing reasons

The concept of a reason is so basic that it is difficult to define it in a noncircular way. What we know for sure is that human beings typically believe certain propositions because they have reasons so to believe, intend to do certain acts because they have reasons so to intend, and so on.

Traditionally, reasons have been classified as either *theoretical* or *practical*: while theoretical reasons are reasons for believing, practical reasons are reasons for desiring, intending, and doing. In turn, practical reasons can be divided into *desire-dependent* and *desire-independent* reasons (Searle, 2001). An obvious example of a desire-dependent reason would be a state of hunger, considered as a reason for eating; typical desire-independent reasons are various forms of obligation, like moral obligations, legal obligations, or the obligations deriving from commitments.

In some cases, we can produce reasons for other subjects making no use of communication; for example, one may set fire to a house to 'convince' somebody to leave the house immediately. To move closer to Grice's definition of meaning, let us now concentrate on a case in which a subject, *x*, provides a reason to another subject, *y*, by making *y* recognize a nonreflexive intention of *x*. Suppose for example that Ann wants to give Bob, a neighbor of hers, a good reason to keep his dog in his yard, instead of letting it run in Ann's garden. Ann may wait for an occasion in which Bob's dog comes in her garden, and then get out of her house with a baseball bat firmly in her hand, making sure that Bob can see her. Unless Bob doesn't care a whit about his dog, he has now a good reason to call it back, even if he does not recognize that Ann intends her intention to beat the dog to be recognized. There are certainly

many kinds of reasons that one can provide to another subject by having a first-level intention recognized, even if the second-level intention to have the first-level intention recognized is not in turn recognized. The problem is to understand what kinds of reasons, if any, require the recognition of a reflexive intention; but before we tackle this issue, we want to discuss a different way of producing an effect in an audience.

Providing affordances

In the dog's example, Ann provides a reason for Bob to call his dog back, counting on his desire to preserve the physical integrity of the dog. While there are certainly many cases in which we try to provide reasons to others, we believe that in many kinds of interaction we do not rely on this.

Consider the following example. On Saturday night, at his favorite bar, Bob spots a charming lady sitting all alone at the counter; he then approaches the lady, and asks her whether she would like to have a drink with him. The effect that Bob is trying to produce in the lady is clear: he would like her to spend some time with him. But Bob is not trying to produce this effect by providing a reason; rather, what he is providing is an *affordance*, that is, an opportunity for action.

Notoriously, the concept of an affordance has first been formulated by James Gibson in the context of his ecological theory of perception (Gibson, 1977, 1979). According to Gibson, an affordance is an objective action possibility offered by an agent's environment. Later on, Donald Norman (1988, 1999) reformulated this concept, defining an affordance as an action possibility that is readily perceivable by an agent. Bob knows nothing about the reasons that the charming lady may have to accept or to reject his invitation; but he knows that she is able to perceive an affordance of the type he is offering her, because drinking something together is common practice in our culture.

As with reasons, to offer an affordance one does not necessarily need to have a reflexive intention recognized. For example, a policeman may disguise as a rich and inexperienced tourist to trap a mugger, and in such a case the policeman obviously wants his or her intention not to be recognized. Again, we have to ask ourselves what sort of affordances can be provided only if one has a reflexive intention recognized.

3.2 The role of reflexive intentions

Grice's definition of meaning is puzzling: on the one hand, the idea that communicating involves the recognition of a reflexive intention nicely corresponds to our intuitions; on the other hand, it is not obvious why such a kind of intentions should be necessary, or even useful. In their book on speech acts, Bach and Harnish (1979) explicitly raise the issue of establishing what kind of effects one may try to achieve through the recognition of a reflexive intention, or R-intention. In their own words (1979:15),

“what sorts of (intended) illocutionary effects—effects consisting in recognition of R-intentions—can there be? In other words, what can be the content of communicative intentions?”

Their answer is that reflexive intentions are necessary and sufficient to “express a thought or, more generally, an attitude, be it a belief, an intention, a desire, or even a feeling.” In the work of Bach and Harnish, “expressing” is a technical term whose meaning is defined as follows (1979:15-16; *S* and *H* respectively denote a speaker and a hearer):

“Expressing: For *S* to express an attitude is for *S* to R-intend the hearer to take *S*’s utterance as reason to think *S* has that attitude.

Accordingly, the intended illocutionary effect (or simply illocutionary intent) is for *H* to recognize that R-intention.”

However, Bach and Harnish do not justify their assumption that the content of a reflexive intention can only be the expression of an attitude: the logical relationship between reflexivity and being a reason to think that a speaker has an attitude is not analyzed.

Let us now elaborate a little more on the dog’s example of the previous subsection. Suppose that Bob sees Ann coming with the baseball bat in her hand, but does not react. As Ann does not really want to beat the dog, when she understands that Bob is not going to call his dog back she crosses the garden and puts the bat in a corner, as if this was what she wanted to do from the start. Nobody can now criticize her for not beating the dog, because at no point of her course of action she publicly committed to having such an intention. This, we believe, is the crucial point: nonreflexive intentions, although they are recognized, do not create commitments. Our hypothesis is that having a reflexive intention recognized is a means to create social commitments, which in turn may constitute reasons or affordances for other subjects.

To understand how social commitments are created, we have to remember two crucial points: first, social commitments are desire-independent reasons for action that are created intentionally; and second, they exist only insofar as they are recognized to exist by the relevant subjects, that is, by their debtors and creditors. In view of these facts, let us see what is needed to create a new commitment.

Let us consider the commitment stemming from a unilateral promise; for example, let us assume that Bob wants to promise to Ann that he will care after her garden while she is in Italy. Following Gilbert’s view of promises, this amounts to Ann and Bob making a joint commitment to the extent that Bob will care after Ann’s garden. As such a commitment can exist only if Ann recognizes it to exist, it has at least to be the case that

(5) Bob intends Ann to recognize that he intends to care after her garden.

Now suppose that Bob does not intend Intention 5 to be recognized by Ann. As satisfying Intention 5 is a necessary step for creating the commitment, this would imply that Bob does not intend Ann to recognize that he intends to commit. But this contradicts our assumption that a commitment may exist only if both the debtor and the creditor recognize it to exist. Therefore, we have to assume that

(6) Bob intends Ann to recognize Intention 5.

By a similar argument we can see that Bob has to intend that also Intention 6 be recognized by Ann. As the argument can be iterated *ad infinitum*, we are led to assume that Bob, in order to commit, has to hold a reflexive intention that can be described as:

(7) (Bob intends Ann to recognize that: he intends to care after her garden and *this*),

where “*this*” refers to the whole expression stated in parentheses.

This condition, however, is still insufficient. To see why, suppose that Intention 7 is satisfied (which means that Ann recognizes Bob’s reflexive intention) and that Bob does not know this. Clearly, such a condition would be insufficient: as a joint commitment exists insofar as it is recognized to exist by both subjects, for Bob to consider himself as committed he must also believe that Ann views him as committed. This leads us to describe Bob’s intention as follows:

(8) (Bob intends it to be common belief of Bob and Ann that: he intends to care after Ann’s garden and *this*),

where again “*this*” refers to the whole expression stated in parentheses.

What we have argued so far is that Intention 8 is necessary to create Bob’s commitment, but we still did not show whether it is also sufficient. Indeed, it isn’t: even if Intention 8 is satisfied, and Bob and Ann attain the relevant common belief, the resulting situation does not yet imply a joint commitment of Bob and Ann. What we still lack is a manifestation of Ann’s intention to play her role in the joint commitment proposed by Bob. It is important to remark that, if Intention 8 is satisfied, Ann is now in the position of accepting or refusing Bob’s *precommitment*, as we call a proposal to jointly commit offered by a subject to another. Independently of the reasons Ann may privately hold to accept, for Ann Bob’s precommitment is an affordance, that is, a readily perceived opportunity for action.

As far as Intention 8 goes, our analysis has been stated in terms of epistemic and volitional states, like common beliefs and intentions. But accepting a precommitment creates a brand new type of entity, more precisely a bundle of deontic relationships between the parties, and we do not think that it is possible to reduce such relationships to epistemic and volitional states alone. A commitment goes beyond this because it builds desire-independent reasons for action.

Let us go back to the issue of accepting a precommitment. When Intention 8 becomes common belief of Ann and Bob, Ann has perceived the affordance of becoming part of the proposed commitment. Let us suppose she decides to accept. What should she do? It seems that she has to make it be common belief of herself and Bob that she intends to become a party of the joint commitment proposed by Bob. For the same reasons we analyzed in connection with Intention 8, Ann’s intention has to be reflexive:

(9) (Ann intends it to be common belief of Ann and Bob that: she intends Bob to care after her garden and *this*).

If Intentions 8 and 9 are satisfied, which means that Ann and Bob enter the relevant common beliefs, their joint commitment will be created.

Intentions 8 and 9 are structurally similar to those required by Grice's definition of meaning, in its reflexive interpretation. The similarity is too striking to be casual. We think there are two possible explanations. The first is that creating a commitment requires expressing intentions with the same level of 'overtness' needed by communication. The second, and more interesting, possibility is that communication presupposes reflexive intentions just because to communicate something *is* to make a commitment, or at least a precommitment. In Section 4 we shall argue in favor of the latter hypothesis.

3.4 Speaker's meaning

We regard the speaker's meaning of an utterance, or more generally of a signaling act, as entirely determined by the speaker's communicative intention, understood as the reflexive intention to produce an effect on a hearer by means of the shared recognition (i.e., common belief) of the intention itself. This far, we do not depart from the Gricean tradition. Where we start departing from it is in our interpretation of the content of a communicative intention.

What kinds of effects can be achieved by means of the shared recognition of a reflexive intention? There are two main points we want to make here. The first point is that, as we argued in Section 3.3, the function of intention reflexivity is to achieve the degree of overtness required by proposing, accepting, refusing, or otherwise dealing with a social commitment of the speaker and the hearer. The second point is that a speaker achieves communicative effects on a hearer not only by providing reasons, but also (and most importantly) by providing affordances, and in particular deontic affordances.

To clarify our standpoint, let us analyze an example. Suppose Albert tells Barbara

(10) "Shall we go to the movies tonight?"

Here Albert expresses his intention to go to the movies with Barbara tonight. By doing so Albert precommits to a joint project of Barbara and himself, and this implies that Albert and Barbara will be jointly committed to go to the movies tonight if Barbara expresses a similar intention (see Section 3.3).

Albert's precommitment need not act as a reason for Barbara to join the project: in an extreme case, it may even be a reason *not* to do so, for example if Barbara is the type of persons who likes to take the initiative herself. In any case, as far as what we want to analyze is the meaning of Utterance 10, being or not being a reason for Barbara's further actions is irrelevant. What is part of the meaning of this utterance is that, if it is correctly understood by Barbara, it provides an affordance for her to commit to a joint project with Albert. Barbara is now in a special position relative to Albert: first, she may take up or not Albert's proposal; second, if she takes it up she may then accept it, refuse it, or deal with it in a more complex way (e.g., by making a counterproposal, or by saying that she would prefer to decide later, and so on).

The previous example shows how we can use the concept of precommitment to analyze the meaning of a proposal to set up a future joint project. Indeed, we believe that a commitment-based approach to meaning allows one to analyze all types of communicative acts. Although a systematic treatment of communicative acts is beyond the scope of this article, we shall now try to show that our working hypothesis is at least promising. To this purpose we now analyze a number of examples of communicative acts of different types.

Promises

Following Gilbert (2006), we regard a *promise* as a joint commitment of two subjects, x and y , to the extent that x will carry out a future action α for y (see Section 2.4). This implies that to be fully in force a promise of x to y must be offered by x and accepted by y . What is commonly called a promise in Speech Act Theory is therefore best viewed as the offer of x to y that they jointly create a valid promise or, in our terminology, as a precommitment of x to y .

It is interesting to compare our view of promises with that submitted by Bach and Harnish (1979:125), who say:

“Commissives are acts of undertaking obligations, but to undertake an obligation is not automatically to create one S 's utterance may express his belief that an obligation is thereby created, but that does not make the belief true even if H shares the belief and it is mutual. That S is obligated to fulfill his commitment is a moral question not answerable by the theory of illocutionary acts.”

If we take “undertaking an obligation” as a synonym of making a precommitment, we agree that doing so is not automatically to create an obligation. However, what it takes to create an obligation is not “a moral question not answerable by the theory of illocutionary acts,” but more simply the acceptance of the speaker's precommitment by the hearer, because the obligations of a promise are not moral obligations, but obligations of a joint commitment.

As an accepted promise is a joint commitment, it entails directed obligation of each subject to the other. Consider for example Albert's promise to Barbara to bring her a bottle of Brunello from Italy. If Barbara agrees, she is then obligated to accept the bottle when Albert brings it to her. Indeed, most promises are only apparently unilateral, because they actually involve some kind of joint project.

Assertions

Speech Act Theory offers two competing views of *assertions*. While according to Bach and Harnish (1979) asserting that p amounts to expressing the speaker's belief that p , in Searle's view (1969:66) it “counts as an undertaking to the effect that p represents an actual state of affairs.” Our view is closer to Searle's, in that we suggest that x asserting to y that p amounts to x making a propositional commitment to y to the extent that p is true.

A propositional commitment of this sort is a social commitment, but at least in general it does not presuppose a joint commitment. This means that assertions create

commitments that do not have to be accepted by the hearer: more precisely, x 's assertion to y that p creates a propositional social commitment of x to p , relative to y , even if y does not in turn commit to p .

Requests

To see how *requests* can be dealt with in terms of commitments, let us first consider a request to perform a future action. A request of this type can be viewed as the mirror image of a promise: by such a request, speaker x proposes to hearer y to create a joint commitment to the extent that y will do a future action for x . A request is therefore a precommitment; if the request is accepted, a joint commitment is created.

It may seem that an 'immediate' request (i.e., a request to perform an action straight away) is different, in that it does not involve creating a joint commitment to a future project. An obvious example of an immediate request is a question, like for example x asking to y

(11) "What is your name?"

As remarked by Clark (1996), however, also this type of communicative acts involves joint projects. By Question 11, x typically precommits to a project of mutual introduction which, if accepted by y , involves directed obligations of both parties: for example, x will be obligated to attend to y 's answer, to memorize y 's name, to tell his or her name back to y , and so on.

There is indeed a difference between a future and an immediate joint project: while a future joint project may be accepted now and performed later, an immediate joint project is accepted by executing it straight away. This means that no future-directed joint commitment is created; rather, the parties' joint commitment is implicit in the very execution of the project. In any case, it seems reasonable to assume that all requests convey a speaker's precommitment to set up a joint project, either future or immediate.

Expressive acts

Searle's *expressives* (Searle, 1976), called *acknowledgements* by Bach and Harnish, include greetings, apologies, congratulations, and the like. All such communicative acts are easily interpreted as parts of immediate joint projects, which are typically common ground and socially relevant. For example, by greeting y subject x carries out his or her part in a mutual greeting project; moreover, x proposes to y to make a joint commitment to the extent that each of them is aware of the other's presence.

Before turning to joint meaning, let us summarize our view of speaker's meaning. First, we regard the speaker's meaning of an utterance, or more generally of a signaling act, as entirely determined by the speaker's communicative intention, understood as the reflexive intention to produce an effect on a hearer by means of the shared recognition (i.e., common belief) of the intention itself. Second, we suggest that reflexive intentions are crucial because every communicative act is intended to create, accept, refuse, or otherwise deal with social commitments, which require the

degree of overtness that only reflexive intentions can produce. Third, we think that the main effect of a communicative act is to provide to the hearer a deontic affordance, that is, the perceived possibility of affecting the network of deontic relationships that bind him or her to the speaker. Of course, much further work is needed to turn these intuitions into a fully fledged commitment-based theory of communicative acts. However, if one accepts at least the general picture, an interesting issue comes to the foreground: as the social commitments of communicative acts involve both the speaker and the hearer, the role of the hearer has to be carefully analyzed. This is what we shall do in the next section.

4. Joint Meaning

The concept of joint meaning that we shall now develop applies to conversational settings; this implies that in certain communicative situations, like in the case of a 'message in a bottle,' there will be speaker's meaning but no joint meaning. Conversations, however, are rightly considered as the basic setting of language use (Clark 1996), and the fact that joint meaning fails to be produced in certain special circumstances should not be regarded as a serious problem.

Conversations are sequences of communicative acts, produced by two or more subjects, each of which has an associated speaker's meaning that depends solely on the speaker's communicative intentions. Moreover, conversations should be regarded as joint activities of two or more subjects, and as such they involve certain types of joint commitments of the participants. In this paper we do not attempt to offer an exhaustive analysis of all types of joint commitments underlying conversations: this matter is very complex, also because different types of conversations are likely to involve different types of joint commitments. Currently we are only interested in a special type of joint commitments, which we call *joint meaning*, whose function is to maintain a shared view of what is said.

Joint meaning is formed every time a speaker and a hearer jointly commit to the fact that a specific communicative act has been performed. Creating joint meaning is therefore a participatory action of a speaker and a hearer: in such a participatory action, the role of the speaker is to perform a communicative act, and the role of the hearer is to take it up. In our view, however, uptake does not only consist in understanding speaker's meaning, or even in achieving common belief that speaker's meaning has been understood. We now try to clarify this point.

4.1 Uptake

We regard uptake as a deontic concept, involving a commitment to the fact that a specific communicative act has been performed. Consider for example the following (obviously unrealistic) exchange:

(11.1) Albert: "Shall we go to the movies?"

(11.2) Barbara: "You just proposed that we go to the movies together."

(11.3) Albert: "I did."

By Utterance 11.1, Albert expresses his intention to go to the movies with Barbara. In doing so, Albert precommits to a joint project of Barbara and himself. Barbara is now in a special position relative to Albert: first, she may take up or not Albert's proposal; second, if she takes it up she may then accept it, refuse it, or deal with it in a more complex way (e.g., by making a counterproposal, or by saying that she would prefer to decide later, and so on).

By Utterance 11.2, Barbara proposes that she and Albert enter a joint commitment to the extent that Albert's utterance was a proposal to go to the movies. Such a proposal is then accepted by Albert (Utterance 11.3). Albert and Barbara are now jointly committed to the fact that Albert proposed to Barbara to go to the movies together: in our terminology, they have achieved *joint meaning* of Utterance 11.1. Of course, this is not how people actually communicate; a more realistic exchange would sound as follows:

Albert: "Shall we go to the movies?"

Barbara: "Great. What about Woody Allen?"

Albert: "Good idea. At Astoria's at seven?"

Here Barbara's uptake is implied by her acceptance of Albert's proposal ("Great") and by her further proposal of choosing Woody Allen's film. Her uptake and further proposal is then taken up by Albert, as implied by his acceptance ("Good idea"). Albert further specifies the joint project by suggesting a place and a time for their appointment. The exchange may then continue with Barbara accepting, refusing, or otherwise dealing with Albert's suggestion.

Communicative acts may be taken up in countless different ways. For example, also a refusal of Albert's initial proposal would imply its uptake, like in the following exchange:

Albert: "Shall we go to the movies?"

Barbara: "Sorry, Albert, I'm too busy"

In actual conversations, there is of course plenty of room for true or simulated misunderstandings, which lead a hearer to take up a different communicative act than has been actually performed: Exchange 2 of Section 2.2 is an example. On the contrary, explicitly refusing uptake seems to be exceptional, even if it is a concrete possibility; suppose for example that George says to Mary,

"The director is a nut."

and consider the following possible answers by Mary:

(12) "I agree."

(13) "I don't agree."

(14) "C'mon, George, you don't really think so."

(15) "I didn't hear you."

By Utterance 12 Mary accepts to jointly commit with George to the judgment that their director is a nut. By Utterance 13 Mary does not jointly commit with George to such a judgment, but she implicitly acknowledges George's unilateral commitment

to the extent that their director is a nut. By Utterance 14, Mary refuses to accept George's unilateral commitment. All these answers imply uptake. On the contrary, by Utterance 15 Mary refuses even to take up George's communicative act. This answer is particularly interesting, because it shows an important feature of a deontic conception of uptake: Mary's answer, "I didn't hear you," implicitly suggests that she did understand George's statement; however, Mary does not take up George's statement, in the sense that she refuses to participate in a joint commitment to the extent that George performed it.

Even if it is a concrete possibility, explicit refusal of uptake appears to be an extreme case in ordinary conversations. Another extreme case is when a hearer limits his or her contribution to pure uptake. This seems to be common in certain special situations, like Rogers-style psychotherapeutic conversations; for example,

Client: "I hate my job."
[silence]

Therapist: "You said you hate your job."

In a normal conversation, uptake is implicit in a hearer's reaction to a communicative act. We shall now analyze the consequences of uptake.

4.2 Achieving joint meaning

We suggest that uptake leads to a form of joint commitment of the speaker and the hearer, more precisely to a propositional joint commitment to the extent that a specific communicative act has been performed. This joint commitment is what we call *joint meaning* of the communicative act.

The joint meaning of a communicative act may or may not coincide with the original speaker's meaning. For example, let us consider again Exchange 2 of Section 2.2. In her first utterance, Evelyn takes up Barry's indirect invitation as a proposal to spend the night working. At this point of the conversation, Barry and Evelyn hold two different views of what is going on. Barry's second utterance, however, implies uptake of Evelyn's construal. Now Barry and Evelyn share the same construal of Barry's initial communicative act: they achieved joint meaning.

It is important to stress that joint meaning is not just common belief of what has been said. Common belief, by itself, does not carry deontic implications: it is possible for two subjects to hold the common belief that p and still act as if they did not believe so. Joint meaning is a joint commitment of two or more subjects, who are then obligated to each other to act coherently; in other words, joint meaning entails directed obligations, rights, and entitlements. Suppose for example that Albert says to Barbara

"I think I'm going for a walk,"

without any intention to invite her to join him, and imagine that Barbara answers:

"Sorry Albert, I'm too busy."

Clearly, Barbara has taken up Albert's statement as an invitation. Albert may now take up Barbara's refusal, thus implying that his original statement was an invitation, by saying for instance

"Pity. Well, I'll be back soon."

At this point of the conversation the joint meaning of Albert's original statement is that it was an invitation. Then suppose that Barbara suddenly changes her mind, and says:

"Oh, well, I think my work can wait. But only a short walk, ok?"

Albert cannot now answer that he actually intended to take a walk on his own, even if this was the actual speaker's meaning of his initial utterance, without violating a joint commitment: by now he has accepted, as a matter of joint meaning, that such an utterance was an invitation, and he is obligated to Barbara to act coherently.

In view of what we said in Section 2.4, the function of joint meaning should be clear: adapting Hamblin's felicitous term, it allows the participants of a conversation to maintain a 'joint commitment store' of past communicative acts, thus creating desire-independent reasons for them to act coherently with what has been said. We are inclined to believe that forming joint meanings is part of what it is to be a conversation; in other words, it may well be *constitutive* of conversations that joint meaning is created and maintained. We regard this point as crucial, also because the very idea that there are constitutive rules of conversation is still a matter of debate (see for example Searle, 1992), but we leave it to future research.

Our concept of joint meaning has several implications, which we plan to analyze in details in future works. A first, obvious implication is that the distinction between speaker's meaning and joint meaning dissolves the apparent conflict between the traditional view of meaning as determined by the speaker's intentions, and Clark's view of meaning as a joint construal of the speaker and the hearer. A further implication is that certain theoretical problems of communication may now be tackled from a different point of view. One issue that may particularly profit from the new perspective is the analysis of strategic ambiguity, which is common in political and diplomatic communication, but also in everyday interactions. Consider for example the case, analyzed by Pinker et al. (2008:833), in which a driver who has just got a ticket indirectly proposes to bribe an officer:

(16) "Gee, officer, is there some way we could take care of the ticket here?"

In a typical situation, after Utterance 16 it will be common belief of the driver and the officer that the driver is proposing an illegal transaction. However, it is still possible for the driver not to enter a joint meaning to such an extent if the officer proves unwilling to be corrupted.

4.3 Speaker's meaning vs. joint meaning

Our declared goal was to reconcile two apparently conflicting approaches to meaning. According to the first approach, which is typical of Speech Act Theory, what a speaker means is basically a function of his or her intentions. According to the second approach (which has been pursued, among many others, by Herbert Clark)

meaning is, in some important sense, jointly construed by the speaker and the hearer. We believe we have reached our goal, by showing that the two approaches aim to clarify two different aspects of meaning, that we respectively call “speaker’s meaning” and “joint meaning,” both of which have a crucial role in ordinary communication.

While speaker’s meaning is solely a function of the speaker’s communicative intentions, joint meaning is a collective construal of the speaker and the hearer. In many cases, the joint meaning of a communicative act coincides with its speaker’s meaning: an intended assertion that p will be jointly construed as an assertion that p , and so on. But at times joint meaning departs, at least to some degree, from speaker’s meaning, whether or not the speaker’s communicative intention has been correctly understood by the hearer. For example, let us consider again Exchange 2 of Section 2.2. Here Barry’s initial communicative act was intended as a pre-request related to an invitation to dinner. By the end of the exchange, the joint meaning of such a communicative act has turned into a pre-request related to a proposal to spend the evening working together. Thus at the end of the exchange Barry is committed to a state of affairs that, in terms of Searle’s analysis of illocutionary acts, is incompatible with the sincerity conditions of his initial communicative act.

On the one hand, it is important to stress that such conflicts between speaker’s and joint meaning imply neither a logical contradiction, nor an immediate violation of the joint commitment that constitutes joint meaning: as we have remarked at the end of Section 2.4, fulfilling or violating a commitment is logically independent of sincerity. On the other hand, when subjects make commitments that are at odds with their personal mental attitudes, they are likely to experience a sense of ‘social pressure,’ because they have now socially based reasons for action that conflict with their personal reasons (see Gilbert, 1996, Chapter 14). For example, going back once again to Exchange 2 we can imagine that Barry initially had a desire-dependent reason to spend a leisure evening with Evelyn. Once the exchange has been completed, Barry has a desire-independent reason to spend the evening working with Evelyn. But probably the initial, desire-dependent reason is still there, and may conflict with the subsequent desire-independent reason.

5. Conclusions and Future Work

In this paper we have proposed to distinguish between speaker’s meaning, understood in terms of personal communicative intention, and joint meaning, understood as a joint construal of the speaker and the hearer. Joint meaning has been defined as a type of propositional joint commitment, more precisely as the joint commitment of a speaker and a hearer to the extent that a specific communicative act has been performed by the speaker. That joint meaning is created and maintained may well be a constitutive rule of conversation.

Our deontic view of communication extends also to speaker’s meaning. We believe that the function of reflexivity of communicative intentions is precisely to propose, accept, or otherwise deal with social commitments of the speaker and the hearer. A

conversation is viewed as a joint activity through which subjects make and maintain a bundle of deontic relationships, producing and exploiting the deontic affordances created by communicative acts. Such deontic affordances may be of two distinct types: those that concern the hearer's contribution to the construction of joint meaning, and those that concern the creation, execution and management of joint projects that may go beyond the sheer production of meaning.

It is important to stress that we do not adhere to the view, indeed common in the recent development of pragmatics, that meaning can be completely explained in terms of a collective construction. The very idea that there may be a conflict between speaker's meaning, understood as a personal mental state, and joint meaning, understood as a collective construction, obviously presupposes that both types of meaning are real and important.

Clearly, our proposal is only at an initial stage, and much work is needed before it can be fully evaluated. More specifically, we plan to develop our research in three main directions. First, we want to gain a better understanding of how deontic affordances are actually achieved and exploited in real conversational situations to produce joint meaning. Second, we feel it is necessary to deepen our comprehension of social commitment, and of joint commitment in particular, as pieces of social reality intentionally constructed by subjects in interaction. Third, it will be important to understand what elements of mental architecture underlie the human ability to form social commitments, and more generally desire-independent reasons for action, which appear to go beyond the epistemic and volitional components of cognition that have been mainly studied so far in Psychology, Philosophy of Mind, and Cognitive Science. To this purpose, we believe that much can be understood from a careful analysis of recent research on the developmental aspects of human sociality and intersubjectivity.

Acknowledgment

We are indebted to Margaret Gilbert for kindly engaging with us in an electronic conversation on joint commitment.

References

- Airenti, Gabriella, Bara, Bruno G., and Colombetti, Marco, 1993. Conversation and behavior games in the pragmatics of dialogue. *Cognitive Science* 17 (2), 197–256.
- Austin, John L., 1962. *How to Do Things with Words*. Harvard University Press, Cambridge (MA).
- Bach, Kent, and Harnish, Robert M., 1979. *Linguistic Communication and Speech Acts*. MIT Press, Cambridge (MA).
- Bratman, Michael E., 1987. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge (MA).
- Carassa, Antonella, Colombetti, Marco, and Morganti, Francesca, 2008. The role of joint commitments in intersubjectivity. In: Morganti, F., Carassa, A., and Riva, G. (Eds.),

- Enacting Intersubjectivity: Cognitive and Social Perspectives to the Study of Interactions*. IOS Press, Amsterdam, pp. 187–201.
- Clark, Herbert H., 1996. *Using Language*, Cambridge University Press, Cambridge.
- Clark, Herbert H., 2006. Social actions, social commitments. In: Enfield, N. J., and Levinson, S. C. (Eds.), *Roots of Human Sociality*. Berg, Oxford and New York.
- Colombetti, Marco, 2000. A commitment-based approach to agent speech acts and conversations. *Proceedings of the Workshop on Agent Languages and Communication Policies, 4th International Conference on Autonomous Agents*, Barcelona, 21–29.
- Gärdenfors, Peter, 2008. Evolutionary and developmental aspects of intersubjectivity. In: Liljenström, H., and Århem, P. (Eds.), *Consciousness Transitions: Phylogenetic, Ontogenetic and Physiological Aspects*. Elsevier, Amsterdam, pp. 281–306.
- Gibson, James J., 1977. The theory of affordances. In: Shaw, R., and Bransford, J. (Eds.), *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*. Lawrence Erlbaum, Hillsdale (NJ), pp. 67–82.
- Gibson, James J., 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.
- Gilbert, Margaret, 1989. *On Social Facts*. Routledge, London and New York.
- Gilbert, Margaret, 1996. *Living together: Rationality, Sociality, and Obligation*. Rowman & Littlefield, Lanham (MD).
- Gilbert, Margaret, 2000. *Sociality and Responsibility: New Essays in Plural Subject Theory*. Rowman & Littlefield, Lanham (MD).
- Gilbert, Margaret, 2006. *A Theory of Political Obligation*. Clarendon Press, Oxford.
- Grice, H. Paul, 1957. Meaning. *The Philosophical Review* 66 (3), 377–388.
- Grice, H. Paul, 1969. Utterer's meaning and intentions. *The Philosophical Review* 78 (2), 147–177.
- Grosz, Barbara J., and Hunsberger, Luke, 2006. The dynamics of intention in collaborative activity. *Cognitive Systems Research* 7, 259–272.
- Hamblin, Charles L., 1970. *Fallacies*. Methuen, London.
- Kuroda, S.-Y., 1989. An explanatory theory of communicative intentions. *Linguistics and Philosophy* 12, 655–681.
- Norman, Donald A., 1988. *The Psychology of Everyday Things*. Basic Books, New York.
- Norman, Donald A., 1999. Affordances, conventions, and design. *Interactions* 6 (3), 38–43.
- Pinker, Steven, Nowak, Martin A., and Lee, James J., 2008. The logic of indirect speech. *Proceedings of the National Academy of Science* 105 (3), 833–838.
- Schiffer, Stephen R., 1972. *Meaning*, Oxford University Press, Oxford.
- Searle, John R., 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge.
- Searle, John R., 1976. A classification of illocutionary acts. *Language in Society* 5, 1–23.
- Searle, John R., 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge University Press, Cambridge.
- Searle, John R., 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press, Cambridge.

- Searle, John R., 1990. Collective intentions and actions. In: Cohen, P. R., Morgan, J., and Pollak, M. E. (Eds.), *Intentions in Communication*. MIT Press, Cambridge (MA), pp. 401–415.
- Searle, John R., 1992. Conversation. In: Parret, H., and Verschueren, J. (Eds.), *(On) Searle on Conversation*. Benjamins, New York, pp. 7–30.
- Searle, John R., 1995. *The Construction of Social Reality*. Free Press, New York.
- Searle, John R., 2001. *Rationality in Action*. MIT Press, Cambridge.
- Searle, John R., 2007. What is language: Some preliminary remarks. In: Tsohatzidis, S. L. (Ed.), *John Searle's Philosophy of Language: Force, Meaning, and Mind*. Cambridge University Press, Cambridge, pp. 15–48.
- Singh, Munindar P., 1998. Agent communication languages: Rethinking the principles. *IEEE Computer* 31 (12), 40–47.
- Strawson, Peter F., 1964. Intention and convention in speech acts. *The Philosophical Review* 73 (4), 439–460.
- Tollefsen, Deborah, 2002. Collective intentionality and the social sciences. *Philosophy of the Social Sciences* 32 (1), 25–50.
- Walton, Douglas N., and Krabbe, Erik C. W., 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany (NY).
- Winograd, Terry, and Flores, Fernando, 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex, Norwood (MA).