Philos Stud (2014) 169:489–508
DOI 10.1007/s11098-013-0205-8

Communication and shared information

Marija Jankovic

Published online: 16 October 2013

Abstract Strawson style counterexamples to Grice's account of communication show that a communicative intention has to be overt. Saying what overtness consists in has proven to be difficult for Gricean accounts. In this paper, I show that a common explanation of overtness, one that construes it in terms of a network of shared beliefs or knowledge, is mistaken. I offer an alternative, collectivist, model of communication. This model takes the utterer's communicative intention to be a we-intention, a kind of intention with a distinctive content that cannot be reduced to an intention in favor of an individual action. I show that the collectivist model can explain overtness in terms of a general feature of we-intentions, namely the requirement that the participants in a shared activity are to intend to act in accordance with meshing subplans.

Keywords Communication · Overtness · Shared intention · We-intention · Common knowledge · Mutual manifestness

In this paper, I argue for (what I will call) the *collectivist* model of communication.[1] This model holds that *standard Gricean* communication is an *essentially intentional*

[1] Note the following about the terminology. [1] The collectivist model of communication acknowledges the existence of the audience's contribution to a communicative action and, therefore, does *not* identify the action of communicating with the utterance act and does not identify communicative intention tout court with the speaker's intention. To align with the prevalent usage, however, talk of an unqualified communicative intention is to be understood as concerning the *utterer's* communicative intention. On the collectivist model, this intention is a we-intention, while on the individualist model, it is an I-intention. (For an explanation of how I am using the terms *we-intention* and *I-intention* see n. 21.) I am explicit

M. Jankovic (✉)
Indiana University, Bloomington, 1033 E. Third St., Sycamore Hall 026,
Bloomington, IN 47405, USA
e-mail: marjanko@umail.iu.edu

collective action type (an EIC).[2] Like line dancing or playing catch, it is a type of action that can only be performed by the utterer and the audience acting together *intentionally*.

By Gricean communication, I mean that type of communicative action on which (1) entails (2).

(1) U communicated that *p*.
(2) U meant something.

On this notion of communication, e.g., Herod's showing the severed head of John the Baptist to Salome does not count as communicating to her that he is dead.[3]

The collectivist model does not hold that *all* Gricean communication is cooperative—lying is an obvious counterexample. But it does hold that cooperation is the *standard* for Gricean communication in the sense that all departures from cooperative communication are to be understood as causally or conceptually dependent on it.[4]

I support the collectivist model by showing that it explains a widely recognized feature of communication—the so-called requirement of *overtness* of communicative intention—in terms of a feature drawn from an independently motivated account of collective intentional action. I contrast this with the failure of a popular family of approaches which construe overtness of communicative intention in terms of the goal of making the intention public (in one sense or another). These, I argue, rule out communication in certain cases in which it clearly occurs and thus fail to provide an adequate account of both communicative intention and overtness.

Section 1 describes the classic example that shows that communicative intention has to be overt. Section 2 shows that the strategy for explicating overtness in terms of various explications of the notion of publicness (namely, the notions of common knowledge and mutual manifestness) is unsuccessful. Section 3 presents the collectivist account of overtness. Section 4 is a short conclusion.

Footnote 1 continued
when I want to talk of the *audience's* communicative intention or a *shared* communicative intention. [2] *Shared*, *collective*, *joint*, *cooperative* are used interchangeably in the context of discussions of shared agency.

[2] For the concept of an essentially intentional collective action type see (Ludwig 2013), especially pp. 4–8.

[3] (Grice 1957) There are other types of information transfer that are called *communicative*. For instance, on an ordinary sense of the word *communicate* the honeybee's dance communicates information about the locations of patches of flowers. But this is not the sense of communication which requires the communicator to mean something by her utterance. Here, our interest is only in the form of communication that entails non-natural meaning.

[4] Jonathan Bennett advocated such a model when he suggested that instead of looking for necessary and sufficient conditions for speaker meaning we secure an explanatory base from which we can "easily range out and capture the rest of the territory" (1976, p. 23). On the cooperative model, the base is cooperative communication. Noncooperative forms of communication are understood as the sabotaging or extending of an established practice.

1 Communication and Strawsonian deception

According to Grice's (1957) analysis, a communicative intention[5] is a transparent intention to get an audience to believe or intend something. More specifically, an utterer U, in uttering x, has a communicative intention just in case:

(G) U utters x intending thereby:

(1) that an audience (A) should produce response r,

(2) that A should recognize that U intended (1),

(3) that A's production of r should be based on A's fulfillment of (2).

A standard objection to (G) is due to Strawson (1964, pp. 446–447). He observes that (G) is compatible with U's intending that A not recognize his intentions (G)-2 and (G)-3. But if U intends this, U does not intend to communicate with A.

Consider a version of Strawson's counterexample to (G) described by (Sperber and Wilson 1996). Mary wants Peter to fix her hairdryer. Not wanting to risk refusal, she does not ask him openly, but leaves the broken pieces lying around as if she were in the process of fixing it. Mary does not intend Peter to be taken in by the staging—she intends him to figure out that the pieces are laid out as a part of her attempt to get him to help. But she intends him to *not* work out that he is intended to reason in this way. Therefore, Mary intends to get Peter to fix her dryer and that Peter recognize this intention. However, Mary does not intend to communicate (in the relevant sense) with Peter. We would not say e.g. that Mary *meant* that Peter is to fix her dryer or that she *told* or *asked* him to fix it. Rather, she got him to fix it by getting him to recognize her intention that he do so.

We can rule out this counterexample by adding (4) to (G).

(4) that A should recognize that U intended (2) and (3)

But the initial counterexample points to a more general problem. (G)-1—(G)-4 do not entail that U intends A to recognize U's intention (4). So it is possible to devise a case in which one has intentions (G)-1—(G)-4 but intends that (G)-4 be hidden from the audience. Likewise for further higher-order intentions. Let us call these cases (in all iterations) *Strawson style cases.* To rule them out, it seems that we have to ascribe to U, for any *n*th order intention that her *(n*-1)th order intention be recognized, an $(n + 1)$th order intention that the *n*th order intention be recognized. But then we ascribe to U an infinite number of intentions. This infinite iteration seems vicious (unlike e.g. an infinite iteration of knowledge states) since each intention requires a distinct mental event that corresponds to its formation.

Grice's preferred response (in 1969) involves describing a structural feature common to all Strawson style cases and adding a postulate to the analysans prohibiting the utterer from having intentions with this structural feature. But this

[5] Grice's target analysandum is utterer's occasion *meaning*. However, Grice is usually understood as holding that an utterer means something by an utterance just in case he *intends to communicate* with some audience, so I take Grice's analysis to apply to communicative intention.

move, as a strategy for replying to Strawson's observation, is (as Grice himself admitted[6]) *ad hoc* and suggests that we lack a satisfying general conception of the cases in which the postulate holds.[7]

2 Common knowledge as the aim of overtness

According to several influential accounts of communication, Strawson style cases reveal that the utterer has to intend not only that the audience recognize the intentions in (G), but that these intentions be—in a specific sense—*shared information* between an utterer and an audience. In this part, I consider (what I will call) the *simple account* and one of its variations—an account by Sperber and Wilson (1996, henceforth S&W), that employs the notion of mutual manifestness. I argue that both accounts are too strong.

2.1 The simple account

The simple account embodies in a straightforward way the motivating idea behind a family of accounts [e.g. S&W's and Schiffer's (1972) approach], namely, the idea that communication involves not just a one-way transfer, but sharing of information.

A proposition that *p* is commonly known in some group G just in case everyone in G knows that *p*, everyone knows that everyone knows that *p*, etc. Let an *informative intention* be [as in Sperber and Wilson (1996, p. 29)] an intention that the audience come to believe something.[8] Then, on the simple account, a communicative intention is (or involves) the utterer's intention that her informative intention be commonly known by her and the audience.

The simple account resolves the difficulty with Strawson style cases. Strawsonian deceivers never intend that their informative intentions be commonly known.[9] But, while attractive, the simple account is too strong. We can show this by considering the so-called coordinated attack problem (CAP).[10]

Two allied armies, commanded by generals G_1 and G_2, are encamped on two hills separated by a valley. The valley is occupied by the enemy. G_1 and G_2 can win easily if both attack at the same time. Otherwise, the attacking army will suffer a

[6] In 1989 (p. 303), Grice says: "the deficiency in that proposal [i.e. adding a clause that prohibits the utterer from having deceptive intentions] was that it gave no explanation of *why* this was a reasonable condition to put into an account of speaker meaning."

[7] Another major type of response in the Gricean tradition is the so-called *reflexive approach* endorsed by Grice in (1957), (Harman 1974, 1977), (Bach and Harnish 1979) among others, and criticized by e.g. (Sperber and Wilson 1996) and (Recanati 1986). I believe that this approach is mistaken because it fails to account for the distinctive commitment to truthfulness an utterer has toward one's addressee as opposed to someone who is openly witnessing a speech act. But I leave a detailed discussion of reflexive accounts to the side for reasons of space.

[8] This assumption restricts our discussion to assertive utterances but the conclusion is general.

[9] For example, in the broken hair-drier case, Mary does not intend that Peter know that Mary knows that Peter knows that Mary has an informative intention.

[10] The description of the case relies on (Chant and Ernst 2008).

catastrophic loss. The generals can try to coordinate an attack by sending messages through a messenger. The messenger has to cross enemy lines, so there is a small chance that he will fail to transmit any particular message. Whenever a message reaches one of the generals, a confirmation of the receipt is sent *automatically* to the other general. This message has the same small probability of not reaching its destination. All of this is commonly known by both generals.

It has been proven[11] that in these circumstances the content of a message sent by one general to the other cannot become common knowledge. Consider the possible outcome in which G_1 sends exactly one message that is received by G_2 but receives no messages. G_1 doesn't know that G_2 has received the message. For rational G_1 will reason as follows. Either **(i)** G_2 did not receive my message or **(ii)** G_2 received my message but G_2's confirmation did not get through. Let ε be the probability that any particular message will fail to reach its recipient. Then the probability of **(i)** given that G_1 received zero messages is $P = 1/(2 - \varepsilon)$ and the probability of **(ii)** is $P_c = (1\text{-}\varepsilon)/(2\text{-}\varepsilon)$.[12] Since $P > P_c$ (for any value of ε), it is more probable that G_2 did not receive the message. Therefore, G_1 does not know that G_2 knows the time of the attack.

Analogous reasoning shows that, no matter how many confirmation messages G_1 and G_2 send to each other, the time of attack never becomes common knowledge between the generals.[13]

Moreover, the fact that G_1 has the *intention to inform* G_2 of the time of the attack cannot become commonly known. The only way for G_2 to come to know about G_1's informative intention is to receive a message; the only way for G_1 to come to know that G_2 knows about his informative intention is to receive a confirmation; and so on. The reasoning above, therefore, applies to G_1's intention to inform G_2 of the time of the attack as well.

Thus, in CAP, an intention that the simple account purports to be a communicative intention, namely, the intention that G_1's informative intention be common knowledge, cannot be satisfied. Given that to communicate is for one's communicative intention to be carried out successfully, the simple account entails that it is *impossible* for the generals to communicate.

But this is implausible. If G_2 receives G_1's message then, regardless of whether her acknowledgement reaches G_1, there is a perfectly intelligible sense in which the generals do successfully communicate (although G_1 doesn't know it). In fact, CAP

[11] By Halpern (1986), Rubinstein (1989).

[12] We calculate P and P_c as follows. G_1 received zero messages, so either **(i)** or **(ii)** obtains. The prior probabilities of **(i)** and **(ii)** are ε and $\varepsilon(1 - \varepsilon)$, respectively. Dividing the credence between **(i)** and **(ii)** so that P: $P_c = \varepsilon$: $\varepsilon(1 - \varepsilon)$ we get that $P = 1/(2 - \varepsilon)$ and $P_c = (1 - \varepsilon)/(2 - \varepsilon)$.

[13] In general, if one of the generals (say, G_2) has received n messages, he knows that there are two possible outcomes consistent with that—either **(i)** G_2's nth message did not get through or **(ii)** it did but G_1's $(n + 1)$th message did not get through. The prior probabilities for **(i)** and **(ii)** are $\varepsilon(1 - \varepsilon)^{n-1}$ and $\varepsilon(1 - \varepsilon)^n$, respectively. This gives us that the probability of **(i)** given G_2's evidence is P and the probability of **(ii)** is P_c.

is often described as a case in which agents who are able to communicate cannot achieve common knowledge of the communicated fact.[14]

Let us consider several challenges to the claim that G_1 and G_2 can successfully communicate in CAP.

Objection 1: We are interested in a form of communication which we have dubbed *Gricean* (see n. 3). But there is an ordinary sense of the word *communicate* (on which e.g. the honeybee's dance communicates information about the location of flower-patches) which does not require the communicator to mean something by her utterance and therefore is not the Gricean sense we're after. So the fact that we can truly say that the generals communicated in CAP does not entail that they communicated in the Gricean sense.

Reply: The mark of Gricean communication is the fact that the intention to engage in it is sufficient for speaker meaning. A paradigmatic communicative act by which utterers mean something is *telling*. That is, if U tells A that *p*, then U means that *p*.[15] And in the example above, when G_1's message reaches G_2, G_1 *has* told G_2 when he intends to attack. Therefore, G_1 has communicated with G_2 in the relevant, Gricean, sense. In addition, G_1 can perform a range of speech acts directed at G_2—warn her that the enemy is on the move, order her to attack at noon, etc.

Objection 2: The foregoing objection to the simple account assumes that communication occurs only if the communicative intention is *satisfied*. But this assumption can be challenged. Drawing from a suggestion by (Strawson 1964, p. 448), it might be said instead that it is only necessary that the communicative intention be *recognized* by the addressee. And this is possible in CAP.

Reply: Let us grant for argument's sake that communicative intention does not have to be satisfied for communication to occur. The simple account still has trouble explaining the *rationality* of forming a communicative intention in CAP. For rational G_1 can realize, before sending the initial message, that his informative intention cannot become commonly known. If so, according to the simple account, he cannot rationally intend to communicate. So the simple account entails that a rational agent in the position of G_1 cannot perform a communicative utterance. But it is implausible that G_1 in CAP cannot rationally intend to produce a communicative utterance. Even if rational G_1 realizes that they cannot coordinate an attack in this way (and that it would be pointless to send a message proposing the time of the attack), it seems possible and appropriate for him to intend to

[14] Here are two examples.

[1] "Making use of epistemic logic it is provable that whenever the smallest uncertainty of message delivery is present, *common knowledge* via *communication* is impossible. (van der Grijn 1994, p. 1, my emphasis).

[2] "[The Coordinated Attack Problem] is a case in which common knowledge is required for action but in which it is impossible for the agents to elevate a piece of information to the status of common knowledge. And this is in spite of the fact that *agents can communicate with each other indefinitely*, have the same interests, and consciously aim to coordinate their beliefs and actions." (Chant and Ernst 2008, p. 553, my emphasis).

[15] Grice's focus in (1957) is explicitly on *telling*. See, for instance, the discussion of Herod's showing Salome the head of St. John the Baptist.

communicate with G_2 about other matters (for instance, to offer encouragement, share intelligence about troop movement in G_2's area, etc.).[16]

2.2 Mutual manifestness

Sperber and Wilson argue that "the appeal to 'mutual knowledge' lacks psychological plausibility" (1996, p. 31). They replace it with an appeal to a weaker notion of publicness, that of *mutual manifestness*.

According to S&W, a proposition (a fact or an assumption, as they put it) is *manifest* to an individual at a given time "if and only if he is capable at that time of representing it mentally and accepting its representation as true or probably true"(p. 39). "To be manifest, then, is to be perceptible or inferable" (*ibid*.).

The notion of manifestness is weaker than that of what is actually known in several respects. First, a manifest proposition need not be true—a proposition is manifest if "the environment provides sufficient evidence for its adoption" (p. 39). Second, a proposition can be manifest to an individual without having been actually entertained by her (*ibid*.). Finally, manifestness admits of degrees: "manifest assumptions that are more likely to be entertained are more manifest" (*ibid*.).[17]

It is *mutually manifest* to A and B that *p* just in case (a) it is manifest to A and B that *p* and (b) it is manifest to A and B that it is manifest to A and B that *p*, etc. For S&W, the infinite regress of manifestness, unlike that of knowledge, is not vicious—no step in this progression has to correspond to an actual state of an individual.

On S&W's account, a communicative intention is the intention to make the informative intention mutually manifest. Mutual manifestness is, according to them, weaker than common knowledge "in just the right way". It is "strong enough to give a precise and interesting content to the notion of overtness"—for though utterers in Strawson type cases intend that their informative intention be manifest, they do not intend that it be mutually manifest. Yet, it is weak enough to circumvent psychological plausibility objections, since "the claim that an assumption is

[16] If we focus on messages about the time of the attack, it seems more plausible that any attempt at communication is not rational.(Chant and Ernst 2008) contains a proof that in CAP common knowledge of the time of the attack is necessary for coordination. No matter how many confirmations reach the generals, once a confirmation inevitably fails to reach one of them, each can reason to the conclusion that it is better to not attack than to attack. If so, it is plausible that a rational G_1 would not even attempt to send a message (for she can calculate that it will be futile). But this does not lend plausibility to the simple account. It only shows that in this case a rational agent can calculate that the (perlocutionary) point of a certain communicative action (i.e. coordination of attack times) cannot be achieved, not that communication is impossible.

[17] In accordance with the gloss S&W give to the notion of degrees of manifestness, I treat it as picking out a purely psychological and not an epistemic (i.e. normative) scale. The notion of manifestness involves epistemic elements, for it involves what the environment provides sufficient *evidence* for accepting. This entails nothing about which propositions are likely to be actually accepted by the individuals in the environment. The notion of *degrees* of manifestness, as I understand it, is a measure of this—that is, a measure of how likely a manifest proposition is to be accepted by actual individuals with thus-and-such cognitive make-up.

mutually manifest is a claim about cognitive environments rather than mental states" (p. 43).

Despite this, S&W's account of communicative intention is too strong. As shown above, the allied generals can successfully communicate in CAP. But G_1's informative intention is not mutually manifest in this environment.

Consider again a case in which G_1 sends one and receives zero messages. G_1 can reason that either **(i)** G_1's message didn't reach G_2 or **(ii)** G_1's message reached G_2 but G_2's confirmation did not get through. This disjunction is entailed by G_1's evidence, and therefore, let us assume, manifest to G_1 (since a proposition is manifest if the environment "provides sufficient evidence for its adoption"). Assume also, as is implicit in the setup, that the only information available to the generals about the (manifestness of the) informative intention is whatever is derivable from the number of the messages they exchange and their knowledge of the setup. Given that, it is manifest to G_1 that if **(i)**, his informative intention is not manifest to G_2.

It is also manifest to G_1 that if **(ii)**, G_2 will receive exactly one message. It will then be entailed by G_2's evidence and, therefore, manifest to G_2 that either **(iii)** G_2's confirmation failed to reach G_1, or **(iv)** G_2's confirmation reached G_1, but G_1's second message did not reach G_2. Assume that the probability that any individual message fails to reach its destination is ε. Then it is manifest to G_1 that if **(ii)**, it is manifest to G_2 that the probability of **(iii)** is $P = 1/(2 - \varepsilon)$ and the probability of **(iv)** is $P_c = (1 - \varepsilon)/(2 - \varepsilon)$ (see n. 12). For any value of ε, $P > 1/2$ and $P_c < 1/2$. Let us assume that a proposition is *not* manifest in an environment E if it is more likely to be false than true given the evidence in E. Then, **(iv)**—that G_2's confirmation reached G_1—is not manifest to G_2. If so (regardless of whether **(iii)** is manifest or not), it is manifest to G_1 that if **(ii)**, it is not manifest to G_2 that G_1's informative intention is manifest to G_2. It follows, by the definition of mutual manifestness, that G_1's informative intention is not mutually manifest.[18]

The argument against S&W's account depends on two assumptions about the relationship of manifestness and evidence, which are, I submit, plausible given S&W's explication of manifestness. They are as follows (letting Δ be the set of propositions that represents the evidence provided by the environment E):

(a) If Δ entails p, then p is manifest in E.
(b) If p is more likely to be false than true given Δ, then p is not manifest in E.

Let us consider possible challenges to each of them.

(a) entails that p is manifest even if the transition from Δ to p involves a prohibitively complex inference or an inference that is not salient to ordinary humans. Therefore, it entails that in any environment there are infinitely many manifest propositions that no one will ever entertain.

[18] It can be proven that no matter how many messages are exchanged the informative intention cannot become mutually manifest in CAP. Note also that while there may be garden variety counterexamples to the simple account (like e.g. a message in a bottle), the formal clarity of the CAP set-up is crucial in providing counterexamples to the account that involves a weaker notion like mutual manifestness.

I believe that such cases are best accommodated in terms of S&W's notion of degrees of manifestness, as ones in which p is manifest to some very low degree. But suppose that we reject (a) and hold that p is manifest in E only if p can be derived from Δ on the basis of an inference that is "psychologically plausible" for the individuals in E. The argument still goes through if *certain* G_1 and G_2 can make and expect each other to make the relevant calculations (even if most ordinary individuals cannot do so). And it seems clear that there can be such individuals—generals who took a class on using probability reasoning in military strategy together, game-theory experts turned generals, etc.

(b) is a rather weak assumption. For example, (b) does not entail that $\sim p$ is manifest. Thus, it does not entail that there is a cutoff probability value v such that, if the probability of p given Δ is above v, p is manifest in E.

One may argue that—since manifestness admits of degrees—we should proportion the degree of manifestness of a proposition in E to its likelihood given the evidence in E. If so, we cannot conclude above that **(iv)** is not manifest to G_2, and the argument is blocked. But such a conception of manifestness is, besides being at odds with S&W's own conception, inherently implausible.

According to S&W, the notion of manifestness is an epistemic notion (for a proposition is manifest if the environment provides *sufficient* evidence for its acceptance).[19] The notion of *degrees* of manifestness is that of a psychological measure of manifestness, i.e. a measure of how likely a manifest proposition is to be entertained by an individual with thus-and-such cognitive make-up (see n. 17).[20]

In addition, on the proposed conception of manifestness, logically inconsistent propositions can be manifest in the same environment. For, suppose that it is manifest that p and the probability of p in E is less than the probability of $\sim p$ in E. If so, $\sim p$ should be manifest as well. But a conception of manifestness that allows this strips manifestness of all epistemic content. A plausible principle concerning evidence is that whatever the evidence available in E makes evident should represent a way the world could be. Given that manifestness is explicated partly in terms of evidence, it follows that the set of propositions made manifest by one's evidence in E should be (at least logically) consistent. It is hard to see how a conception of manifestness that fails to meet this constraint could preserve a link between manifestness and evidence.

Finally, suppose that we allow that manifestness does not have to satisfy epistemically motivated constraints [even as weak as (b)] and consider how such a conception would play out in CAP. If we allow that both **(iii)** (G_2's confirmation failed to reach G_1) and **(iv)** (G_2's confirmation reached G_1, but G_1's second message

[19] I read "sufficient" in S&W's "the environment provides evidence *sufficient* for adoption" of a proposition p normatively throughout. It can perhaps be given a non-normative, psychological reading, but I explain below why a nonepistemic reading is implausible.

[20] Some remarks S&W make support (b) as the correct interpretative assumption. S&W consider a case in which (it is manifest to Peter that it is manifest to Mary)^n that the phone is ringing. They say that, even though more complex propositions in the progression are not likely to be entertained, they are manifest since "there is no cut-off point at which these assumptions are more likely to be false rather than true; they remain manifest throughout, even though their degree of manifestness tends asymptotically toward zero" (p. 42, my emphasis).

did not reach G_2) are manifest to some degree in G_2's environment, we may be able to maintain that G_1's informative intention is mutually manifest. But, **(iii)** is manifest as well, ensuring that G_1's informative intention is also not mutually manifest (and that this is manifest to G_1). S&W describe successful communication as "extending the mutual cognitive environment" (p. 64) of the utterer and the audience. But, on the present conception of manifestness, G_1's utterance, even though it is successful, fails to do so. For G_1's informative intention both is and isn't a part of G_1's and G_2's mutual cognitive environment, and this is manifest to G_1 and G_2. This conception of manifestness, then, is also incompatible with S&W's overall view of the point of communication.

Our discussion shows that the picture of successful communication as necessarily resulting in shared information is mistaken. This is revealed in contexts like CAP in which communication is successful despite the fact that the intention with which an utterance is produced is not, in the relevant sense, public. This leaves us with the problem of explaining both the nature of the requirement of overtness and its source, to which I turn in Sect. 3.

3 Overtness and shared agency[21]

On the collectivist model of communication, standard communication—communication between sincere utterers and trusting audiences with the aim of transferring information—is an EIC (see p. 1). In this part, I show that the collectivist account provides an independently motivated explanation of overtness that is immune to counterexamples.

First, I describe several features of shared intentional agency important for the subsequent discussion. Second, I briefly motivate the collectivist account of communication. Third, I show why an individual's intention in favor of a joint intentional activity has to be overt. Finally, I reply to some objections.

3.1 Shared agency and intentions

Collective intentional action is not a mere aggregation of individual actions. To see this compare a group of people in the park running toward a centrally located shelter to escape the rain with the same group performing the same movements as a part of a ballet troupe's outdoor performance.[22] A collective intentional action is performed only in the latter case, but the difference between the cases does not lie in the

[21] Note the following about the terminology in this part. I take a *we-intention* to be an *intention* an individual agent has when participating in a collective intentional action [pace (Searle 1990) who uses *we-intention* for a *type* of mental state distinct from individual intention]. A we-intention is a state of an individual and it is not to be confused with a *shared intention*, which is a state of a *group* that intends to do something. I will call *I-intentions* intentions of individuals that are not *we-intentions*. I-intentions can be directed at individual or collective actions (as in the example of my I-intention that we run out of the theatre).

[22] The example is from (Searle 1990).

individual behavior, since it is the same in both cases.[23] I assume, in line with one of the main positions in the current discussion of shared agency, that the difference between a mere aggregation of individual actions and a joint intentional action lies in the content of the participants' attitudes, specifically, the intentions with which they act.[24]

If we perform or intend to perform a collective action *J* intentionally, we may be said to *share* an intention to *J*. In such cases it may also be said that *we* have an intention to *J* or that a *group* intends to *J*. We will understand shared intention as a structure of interrelated attitudes of individuals, not as a single attitude in the mind of a "group agent". A group "intends something" or its members "share an intention" only if each has the appropriate *we-intentions* (and, perhaps, additional attitudes).[25]

We-intention is distinct from **(i)** an intention to perform one's part in a collective action, even in circumstances of common belief that each member will do her part, and **(ii)** an intention in favor of a collective action (or in favor of bringing about an outcome collectively).[26] In support of **(i)**, consider an example from (Searle 1990). Each member of a group of business school graduates, convinced of Adam Smith's theory of the invisible hand, resolves to do his part in helping humanity by acting as selfishly as possible without regard to anyone else. It is a common belief among the graduates that each intends to do this. Even so, they do not intend to help humanity together. In support of **(ii)**, consider that I may intend to bring it about that we run out of the theater by yelling "Fire!" This is not a we-intention and even if each of us had such an intention (and even if this was common knowledge), we would not share an intention to run out of the theater.

Following Bratman,[27] I assume that our *shared intention* that we J is realized in part in the following structure of intentions of individuals.

(1) (a) I intend that we J and (b) you intend that we J.
(2) I intend that we J in accordance with and because of 1*a*, 1*b*, and meshing subplans of 1*a* and 1*b*; you intend the same.[28]

[23] Note that collective action does not have to be intentional under any description (as e.g. when we pollute the environment together but not intentionally). Agents in the first case perform a collective action of e.g. running together, but not a collective intentional action.

[24] For arguments see e.g. Bratman (2009), Ludwig (2007).

[25] The assumption that a group of agents shares an intention to J just in case each member of the group has a we-intention that they J may be questioned in the case of institutional groups. But this complication is tangential to our concern here, as we will be concerned with the agency of informal, non-institutional groups, for example, the ones we form when we have a conversation, walk together, etc.

[26] There is wide agreement on both of these points. See e.g. Searle (1990), Bratman (1993), Gilbert (1990).

[27] See (Bratman 1993, 1999b, 2009). Bratman's account is the most well-known of this kind. But the discussion that follows can be stated in a number of other frameworks.

[28] Bratman also requires that (1) and (2) be common knowledge. I discuss why I omit this requirement below. Bratman takes the conditions in his analysis of shared intention to be jointly sufficient but allows that shared intention may be multiply realizable. See e.g. Bratman (1999a, pp. 143–144; 2009, p. 155).

That is, if we share an intention to J, then each of us has an intention in favor of our J-ing. These intentions *interlock* in the sense that each intends the joint activity to go in part by way of the other's intention in favor of the joint activity. Finally, we should intend that the sub-plans associated with our intentions *mesh*—i.e. that they be mutually consistent. Bratman's formulation focuses on shared intention. We will take an individual we-intention to (at least) involve an intention that we J in accordance with and because of this intention, your intention that we J and meshing subplans of these intentions.

3.2 The collectivist model[29]

Given this background, the main claim of the collectivist model can be formulated as follows. Communication is an EIC, and individual *communicative intentions* are we-intentions. They come in two varieties (characterized by the two distinctive roles in the activity)—a we-intention of an utterer to perform a communicative action and a corresponding we-intention of the audience. The goal of (standard) communication is, as Grice observed, that the audience have a certain response (e.g. a belief). But this goal is something that the utterer and the audience aim at together intentionally, not something that the utterer I-intends to produce in an audience.

The collectivist model of communication is on the face of it a plausible account of our ordinary communicative exchanges. The minimal unit of a paradigmatic communicative exchange (at least typically) involves cooperation between two participants where the role of one of the participants is to (e.g.) speak to the other and of the other to cooperatively attend to what the speaker says. (That attending is an action is clear from the fact that eavesdropping is an action: the eavesdropper *does* what the audience of a communicative act is *intended to do*.)[30]

A tacit assumption of Gricean accounts considered above is that a communicative intention is an I-intention. Depending on whether the audience is conceived as an agent of the "production of response *r*", the intentions deemed communicative by the accounts discussed in Sects. 1 and 2 can be seen as directed at an individual or at a collective action.[31] But on either reading they are not intentions in favor of a collective intentional action with the audience.

[29] This is a minimalist sketch of the collectivist model. Further development must wait another occasion. The sketch, however, suffices for present purposes. It is developed enough to show how it solves an important problem and this, in turn, motivates a fuller development of it.

[30] It might be suggested that communication can be effective even without cooperation on the part of the audience, if the utterer succeeds in capturing the audience's attention without her will. But relying exclusively on attention capture to communicate would be difficult and inefficient, so this can hardly be a paradigmatic case of communication or a suggestion for how it is to be conceived as a going enterprise. I would suggest that this case is like a case in which you get someone to catch a ball by throwing it at him unexpectedly. This doesn't show that playing catch isn't a collective intentional activity. *Mutatus mutandis* for communication. In this paper, however, I put aside further discussion of this issue. My goal here is to show how the collectivist model helps us understand overtness of communicative intention. I leave for another occasion responding to objections of the sort just sketched (where we move away from paradigmatic communicative exchanges to other sorts that those make possible).

[31] The former is the case if the production of the response is caused solely by the utterer's action.

If this is right, then these accounts cannot capture something essential about communicative action (in particular, about the content of communicative intention). From the perspective of the collectivist account, the irreducibility of shared intention to agglomerations of I-intentions is what explains why attempts to modify Grice's analysis in order to accommodate Strawson style cases have had an air of a "wild clause-chase", as they set off after an impossible kind of reduction of an inherently social phenomenon.

3.3 The collectivist model and Strawsonian deception

The collectivist model affords an account of overtness of communicative intention in terms of a general feature of we-intentions. The explanation, in short, is this: an intention to get an audience to have a response by engaging in Strawsonian deception is not a communicative intention because the audience is not intended to have the response as a result of acting together intentionally with the utterer.

Strawsonian deception seems to be in general incompatible with joint intentional action. Examples of deception modeled after Strawson style cases are intuitively not examples of intended cooperation. Consider the following case. Anne intends that she and Bob paint their houses the same color, green. She intends to bring this about by pouring green paint into a bucket left under Bob's balcony when she knows that Bob will be on the balcony looking down (but when Bob doesn't know that Anne knows that he is looking). Anne intends Bob to think that she is secretly leaving paint at his door and to reason as follows: "The paint she is leaving is green, the same color she intends to paint her house, so it must mean a lot to her that we paint our houses the same color. I don't care about the color of my house that much, so I will take pity on her and do it."

If Anne and Bob came to intend to paint their houses the same color as a result of Anne's plan, we would not say that they shared an intention to do it. Rather, Anne took advantage of Bob's helpful disposition to get what she wanted.[32] This lends support to the idea that the incompatibility of communicative intention with Strawsonian deception is to be explained in terms of a general feature of collective intentional action.

What is this general feature? In talking about Strawsonian deception as a type, we are recognizing that the cases (in all iterations) have something in common. The following feature stands out. In each case one agent, U, has a plan, P_U, that involves getting another agent, A, to do or bring about something, φ. (φ can be A's part in a collective action.) P_U involves getting A to (intend to[33]) φ in accordance with a

[32] Too see this, note that the complex of intentions arrived at through Anne's deception cannot perform the characteristic roles of shared intention. For example, perceived conflict or problems in carrying out the joint activity would not trigger the type of shared deliberation characteristic of shared agency. It could only trigger furtive sub-plan adjustment. Thus, if Bob runs out of paint, he can ask Anne for more only under false pretenses (given that his intention is that she *not* realize that he was looking at her leaving the paint at his door).

[33] I assume, in accordance with what Bratman (1984) calls the *Simple View* that if A intentionally φ-s, then A intends to φ. For a survey of the debate about this issue see Mele (1992).

certain plan—call it P_A—such that: **(1)** A believes that it is *not* the case that U intends her to φ in accordance with P_A and **(2)** it is part of A's total plan that U is not to come to know that A φ-ed in accordance with P_A.[34] Call this feature of P_U F_{GEN}.

Anne's plan for bringing it about that she and Bob paint their houses green has F_{GEN}. She intends (P_U) that Bob form and execute the intention (P_A) to paint his house using the paint *she* left on his porch, that **(1)** he believe that Anne does not intend him to paint in accordance with P_A, and that **(2)** Bob work so as to keep it hidden from Anne that this (i.e. P_A) is the plan he intends to act in accordance with.

Same holds of classic Strawson style cases. In the broken hairdryer case, Mary intends (P_U) that Peter come to believe that she wants him to fix the dryer by virtue of (P_A) attending to Mary's arranging of the broken pieces *as if she were trying to fix it*.[35] If Mary is successful, **(1)** Peter believes that Mary does not intend him to acquire the belief in accordance with P_A (he thinks that she intends him to not think that she left the parts lying around *as if* she were trying to fix the dryer herself), and **(2)** Peter intends to keep from Mary that he intends to act in accordance with P_A.

If U intends to act and get A to act in accordance with a plan that has F_{GEN}, then it is not the case that U intends that he and A act in accordance with and because of meshing subplans, and therefore not the case that U intends to act together intentionally with A. U knows that if his intention is satisfied, A intends to act so that U does *not* come to know that A intends to φ in accordance with P_A. So A's intention is satisfied just in case U does not come to know that his deceptive plan was successful. Perhaps in the standard deception cases the deceiver does not typically intend specifically to find out whether his intention is satisfied (he may merely hope that it is). But the deceiver surely is not committed to being kept in the dark about it.[36] Insofar as U is open to finding out whether his intention—that A φ

[34] Grice (1969) formulates the characteristic feature of Strawsonian deception as follows. In each case there is an inference-element E such that U uttered x intending both that (i) the inference by which A reaches his response r should rely on E and (ii) that A should think U to intend that (i) be false. The description of the characteristic feature above differs from Grice's in that it describes the deceptive plan as requiring not only that A should think that U intends that (i) be false but that A *intends* that U believe that (i) is false (i.e. that A intends that U believe that her deception, as U conceives of it, is successful). To see that this additional requirement is apt consider again S&W's Strawson style example. Let us say that *p* is the proposition that Mary wants Peter to fix the drier. Mary's deceptive plan is carried out just in case Peter, believing that *p* and believing that Mary intends that he does not believe that *p*, intends that Mary maintain the false belief that her plan was successful. If Peter lets Mary know that he has seen her lay out the parts and that he recognizes why she did it, Mary's deceptive plan does not succeed. See, however, n. 36 for an explanation of how the collectivist account can accommodate deception cases in which U lacks the intention that A should *intend* that U believe that (i) is false.

[35] It might be objected that Peter does not act in accordance with any *plan* in coming to believe that Mary wants him to fix the drier. He does not do any prior planning but simply finds himself presented with scattered parts. This is not an objection on the present use of "plan". I understand intention as an attitude of a commitment to a plan of action. "Plan" here means simply a sequence of action types. A plan is something that can be a product of explicit planning, but need not be.

[36] This discussion depends on treating (2) as a necessary feature of Strawson style deception. This was crucial in explaining why U does not intend that he and A act in accordance with meshing subplans, since U does not intend that the intention A has in (2) be satisfied. One may wonder whether a deceptive intention that included (1) but not (2) would be a we-intention (and a communicative intention). Consider

in accordance with P_A—is satisfied, U is not committed to the activity proceeding by way of meshing subplans.[37]

Note that it is not that U intends A to act in accordance with a subplan that is not co-realizable with his. P_U and P_A can both be realized, for example, if U never finds out that his deceit is successful. But U does not have a positive intention in favor of achieving this mesh—he is open to "somehow [sidestepping], possibly using deception, the sub-plans of the other" (Bratman 2009, p. 158). This undermines the collective intentionality of their action. (Consider this case as perhaps a clearer analogy. Suppose that U and A intend to paint the house together. A intends to buy the brushes, and U intends to buy the paint, A intends to scrape and U intends to paint. U does not intend to buy the brushes, but is open to doing it and perhaps will do so if the paint store has some. U also does not intend to scrape, but may change his mind in the morning. U has no intention to negotiate these changes with A, and the fact that these changes would make his plans inconsistent with A's is not a consideration against them. I do not think that in this case we would say that U and A share an intention to paint the house together.)

Therefore, an intention to act in accordance with a plan with F_{GEN} falls short of a we-intention, because the deceiver does not intend that the joint activity come about in accordance with and because of meshing subplans.[38]

The upshot is this. On the collectivist model, the incompatibility of communicative intention with Strawsonian deception is traceable to the intention being directed at a joint intentional activity of speaker and audience, in which the speaker's subplans are to mesh with the audience's subplans directed at their

Footnote 36 continued
the dryer case again. Suppose that Mary conspicuously arranges the parts, that Peter reasons in the way she intends him to, and that he as a result comes to believe that she wants him to fix the dryer. But Peter does not intend to keep it hidden from Mary that he saw her arranging the parts. He intends to tell her that he saw her and that he is happy to fix her dryer. But he simply doesn't get to. Suppose that this state of affairs satisfies Mary's deceptive intent. In such a case, I do not think that we should say that Mary's intention in arranging the parts was a communicative one or that Mary and Peter acted together intentionally in getting Peter to have the belief in question. But, we cannot invoke the explanation above, as it is not the case that Mary is open to sidestepping any of Peter's intentions. However, here Mary intends that Peter have a false belief about how the mesh is to be achieved. This suggests that parties that share an intention to J need not only to intend to J in accordance with meshing subplans but also that no party should intend that the others have a significant misconception about how she intends to achieve the mesh. (What counts as "significant" misconception may be a matter of degree. A certain amount of sneakiness may not undermine a shared intention.) Requiring something along the lines of what is proposed in (Kutz 2000, p. 6)—that the participants ought to have "dispositions favorable to mutual manifestness"—would suffice to rule out shared intention in these deception cases. Thanks to the anonymous referee for raising this issue.

[37] My thanks to Michael Bratman and Kirk Ludwig for help with this.

[38] This discussion treats the intended mesh as a *necessary* component of a shared intention. Bratman, however, claims that his conditions are merely jointly sufficient. (See n. 28.) In that case, he says, "the indicated web of interlocking intentions was one important case of shared intention" (1999a, p. 144). If intended mesh is not necessary for shared intention, this discussion is limited to the important case of shared intention identified by Bratman and exemplified by typical cases of e.g. intending to go to NYC together, take a walk together, sing a duet, etc. of which communication seems to be a paradigmatic instance.

communicating. In particular, Strawsonian deception is incompatible with the intention that their subplans mesh. Thus, we see how the collectivist model, in explaining what goes wrong in cases of Strawsonian deception, explains why communication requires overtness without requiring the publicness of the communicative intention.

3.4 Objections and replies

Let us consider two potential objections to the collectivist model.

Objection 1: People often enter into communicative exchanges spontaneously, without a prior agreement to do so. For example, the audience of a communicative action may simply find herself being addressed by the utterer. It is implausible that in such cases the utterer and the audience share an intention to communicate.

Reply: Joint intentional action often happens as a result of a *prior* shared intention. One of the main roles of a shared intention is to coordinate interpersonal planning leading up to the action. But prior shared intention and associated shared planning are not necessary for joint intentional acting. This is clear if we consider EIC-s, especially those that embed explicit constitutive rules. Consider playing tennis. The parties need not engage in any shared planning about how to play—a collective action plan with roles for each participant is simply a part of the knowledge of what it is to play tennis. Games of tennis are typically initiated by prior agreements, but this too is absent from a lot of our shared activities. Even a game of tennis may be initiated by a player's serving a ball toward an unexpecting opponent. And similarly there need be no prior planning when one takes part in a stadium wave, starts an applause, offers money to a cashier, or initiates a conversation.

Objection 2: (CK) [or, alternatively (MM)] expresses a necessary condition of sharing an intention:

(CK) If we intend to J, then it is common knowledge that I intend that we J and that you intend that we J.[39]

(MM) If we intend to J, it is mutually manifest that I intend that we J and that you intend that we J

If so, CAP presents a problem for the collectivist account as well, since it is impossible to achieve common knowledge/mutual manifestness of the relevant intention in CAP and, therefore, impossible to perform a joint intentional action.

Reply: While (CK) and (MM) look plausible when we focus on shared activities that require responsiveness in planning and execution (such as singing a duet or painting a house), they are less so when we consider the full range of shared intentional behavior. As Kutz (2000, p. 1) notes, shared intentional activities range "from intricate duets to routinized, hierarchical cooperation within bureaucratic structures". Given the diversity of shared agency, it is implausible that there are universal epistemic conditions for shared intention.[40]

[39] See e.g. Bratman (1993, 2009), Tuomela (1991).

[40] See Chant and Ernst (2008).

Here we only need a more modest claim. Suppose that J_{12}-ing is an activity with the following features: (1) each participant knows what her part consists in (either because J_{12}-ing is fully planned beforehand, or because the concept of J_{12}-ing is the concept of an EIC with determined roles for each participant, or because the participants can determine what to do without shared planning), and (2) at least one participant can make her contribution independently of the other participants. The modest claim is this: a group can share an intention to J_{12} without knowing and without it being manifest to the members that others also intend to J_{12}. It is sufficient that each believes that it is *possible* that the others intend (and carry out the intention) to J_{12}.

Consider this example. X is a member of a team of spies with a mission to gather information about the enemy's nuclear program. The mission is pre-planned extensively and members of the team are told to have no contact with each other. X has reason to believe that certain members of the team are double agents and do not intend to do their parts. But convinced of the utmost importance of the mission, X does her part, hoping (perhaps foolishly) that the other members will do their parts as well. X does not know nor is it manifest to her that the other members intend to J. But if every member of the spy team ends up doing her part (i.e. X's suspicions are in fact incorrect), it is reasonable to say that they carried out the mission together intentionally. The same goes if each member of the team shares X's suspicions but does her part nevertheless.

Some activities with features (1) and (2) can involve a certain level of responsiveness in execution. These are usually activities where one member acts as an initiator and hopes that the other, recognizing her intention and the plan behind it, follows through on it. Consider this example from (Kutz 2000, p. 18). It starts to rain while X and Y are having a picnic. X grabs the sandwiches and runs toward the car. X hopes that Y will grab the drinks and the blanket (as Y's part in their saving the picnic). If Y does, then X and Y acted together intentionally in saving the picnic. But at the time of acting X does not know that Y intends or will come to intend to save the picnic. It may easily be true that X had no reason to believe that Y will come to so intend—for example, X may think that Y is given to dozing off. Thus, it is not manifest to X that Y intends or will come to intend to save the picnic. As Kutz notes, X may initiate their J-ing having nothing more than a weak hope that Y will follow suit. In these activities Y's contribution is responsive to X's, but even so, the requirement that Y know or that it be manifest to her that X intends to J is too strong. Y can respond to X's contribution simply hoping that X intended the contribution as his part in their J-ing.[41]

It follows that, if J-ing is an activity with features (1) and (2), common knowledge/mutual manifestness that everyone intends to J is not necessary for

[41] Note that, with shared intentional activities in which the participants' contributions are not simultaneous, the initiator doesn't have to believe that it is possible that other members intend to J at the time of his contribution to J-ing. It is enough that he thinks it is possible that they will come to intend to J at the time at which they are to make a contribution. Thus, we can sing a song together in virtue of my singing the first part simply hoping that you will catch on in time to sing the second part.

shared intention. Communicative action has both of these features. It is an EIC and what it is to do a part in it is determined by the conception of the action type itself. And an utterer can initiate a communicative action without knowing or it being manifest to her that the audience intends to communicate as well.

Can the collectivist account, therefore, explain the possibility of communication in CAP?

The following pair of intentions can be satisfied in CAP.[42]

(i_1) G_1 intends that G_1 and G_2 bring it about that G_2 acquire a belief about when G_1 plans to attack in accordance with and because of this intention, G_2's corresponding intention, and meshing subplans of these intentions.

(i_2) G_2 intends the same.[43]

This pair of intentions is satisfied if, as a result of (i_1), G_1 sends the messenger with the message that he plans to attack at, say, noon, and G_2 as a result of (i_2) attends to the message and comes to believe that G_1 plans to attack at noon.

Agents in the CAP can also satisfy the epistemic conditions for shared intentional action outlined above. Rational G_1 will believe that it is possible that G_2 has or will come to have the intention (i_2). The likelihood of the message getting through is high and G_2 has an interest in finding out its content. And G_2 will believe that it is possible (in typical circumstances) that G_1 has the intention (i_1) upon encountering the messenger. Therefore, the collectivist model shows how communication is possible in CAP.

It may be suggested that the iteration we see with common knowledge will be expressed with the weaker epistemic condition as well. For example, it may be argued that if X and Y both believe that it is possible that the other we-intends to J,

[42] Actual communicative intentions are more complex because communication is not just any (joint intentional) way of imparting a belief. For example, the appropriate way will restrict the utterances to actions that are not natural signs of the belief content. But (i_1) and (i_2) serve our purpose and simplify the discussion as there is no reason to think that introducing the relevant restrictions will make it impossible to impart a belief in such a way in CAP.

[43] Bratman's account of shared intention entails that if U and A share an intention to J, each intends that they J. It follows that, when there is no prearranged topic, A and U cannot share an intention to *communicate that p* (for some specific *p*). The audience cannot intend that they communicate that *p* when she does not know in advance what signal the utterer is going to produce. (Thanks to Karen Lewis for bring this up in her prepared comments at the 2012 Central APA meeting.) We bracket this concern here and assume that the generals have a time of the attack as a prearranged topic in stating the content of their intentions. This does not affect the overall point. I maintain that in cases with no prearranged topic the utterer and the audience share an intention to *communicate something or other*, that is, to engage in a determinate of the determinable type *communicate*. The formal implementation of this idea does not fall in the scope of this paper. Informally, we consider a typical communicative intention to be shared at the level of a determinable, partial plan. The utterer fills in this partial plan by making a specific utterance and that determines the audience's part. This is exactly analogous to e.g. playing tic-tac-toe or chess (and, more generally, to engaging in EICs in which participants take turns and in which the contributions at each turn constrain subsequent contributions), where the joint intention (prior to the first move) is to play a game of the type characterized by the rules of tic-tac-toe and where each move constrains the choices of the subsequent moves for each player.

but it is not the case that X believes that it is possible that Y believes that it is possible that X we-intends to J, there is no shared intention. I will not discuss here whether the iteration is necessary.[44] But it is worth noting that there is no reason to think that such conditions cannot be satisfied in CAP. In CAP, from the fact that G_1 received exactly *n* messages it follows that it is not the case that [G_1 knows that G_2 knows]$^{n+1}$ that G_1 intends (i_1) (and *mutatis mutandis* for G_2). But from the fact that G_1 received exactly *n* messages it does not follow that it is not the case that [G_1 knows that it is possible that G_2 knows that it is possible]$^{n+1}$ that G_1 has (i_1). Consider a case in which G_1 sends one and receives zero messages. He thinks that it is more likely that G_2 did not receive his message. From this G_1 can conclude that it is not the case that G_2 knows G_1 has (i_1). But G_1 cannot conclude that it is not possible that G_2 knows that it is possible that G_1 has (i_1). In fact, rational G_1 should think that, even though it is more likely that his message did not get through, it is not impossible that it did—it is only slightly less likely for small values of ε. Therefore, weaker iterated conditions can be satisfied in CAP.

4 Conclusion

In this paper, I have shown that a common Gricean explanation of overtness of communicative intention, one that construes it in terms of a network of shared beliefs or knowledge (actual or potential) of the utterer and the audience, is mistaken.

I offered an alternative—collectivist—model of communication that rejects the standard Gricean assumption that the speaker's intention is an I-intention. The collectivist model takes it to be a we-intention, a type of intention with a distinctive content. From the perspective of the collectivist account, an initial symptom of the limitations of the individualist conception were the Strawson style counterexamples. They showed that there was a sense, not captured by the standard formulations of the Gricean mechanism, in which the intention with which the speaker acts must be open or transparent. Explanations of this phenomenon within the individualist framework had an ad hoc character, as Grice himself was aware.

The collectivist framework provides a solution to this problem. It works by exhibiting communicative action as falling into a broader category—that of essentially intentional joint activities. The requirement of overtness is explained as a necessary condition on participatory intentions. This gives us a new, and better, purchase on communicative intentions, and holds out the hope of making fresh progress on the project of showing how the right structure of intentions can be a source of linguistic meaning.

[44] Note, however, that for Gricean accounts the main motivation for iterated epistemic conditions came from Strawson style cases. With the collectivist model there is no parallel motivation—we can rule out Strawsonian deception without introducing any epistemic conditions on the utterer's intention. It is of course still open that the iterated conditions are a general requirement on a shared intention.

Acknowledgments I would like to thank Michael Bratman, Gary Ebbs, Mark Kaplan, Karen Lewis, Matthew Stone, Steven Wagner, Joan Weiner, and an anonymous referee for this journal for helpful comments. I owe special thanks to Kirk Ludwig for discussions and written comments on multiple versions of this paper. I developed large parts of the critical portion of this paper at the International Summer School in Cognitive Sciences and Semantics in Riga, July 2012. My work there benefited from various conversations with the participants.

References

Bach, K., & Harnish, R. (1979). *Linguistic communication and speech acts*. Cambridge: MIT Press.
Bennett, J. F. (1976). *Linguistic behaviour*. Cambridge: Cambridge University Press.
Bratman, M. E. (1984). Two faces of intention. *Philosophical Review, 93*(3), 375–405.
Bratman, M. E. (1993). Shared intention. *Ethics, 104*(1), 97–113.
Bratman, M. E. (1999a). I intend that we J. In *Faces of intention* (pp. 142–163). Cambridge: Cambridge University Press.
Bratman, M. E. (1999b). Shared cooperative activity. In *Faces of intention* (pp. 93–108). Cambridge: Cambridge University Press.
Bratman, M. E. (2009). Modest sociality and the distinctiveness of intention. *Philosophical Studies, 144*(1), 149–165. doi:10.1007/s11098-009-9375-9.
Chant, S. R., & Ernst, Z. (2008). Epistemic conditions for collective action. *Mind, 117*(467), 549–573. doi:10.1093/mind/fzn033.
Gilbert, M. P. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy, 15*(1), 1–14.
Grice, H. P. (1957). Meaning. *The Philosophical Review, 66*(3), 377–388. doi:10.2307/2182440.
Grice, H. P. (1969). Utterer's meaning and intention. *Philosophical Review, 78*(2), 147–177.
Grice, H. P. (1989). Meaning revisited. In P. Grice (Ed.), *Studies in the way of words* (pp. 283–303). Cambridge: Harvard University Press.
Halpern, J. Y. (1986). Reasoning about knowledge: an overview. In *Proceedings of the 1986 conference on theoretical aspects of reasoning about knowledge* (pp. 1–17).
Harman, G. (1974). Review of Meaning by Stephen R. Schiffer. *The Journal of Philosophy, 71*(7), 224–229. doi:10.2307/2025349.
Harman, G. (1977). Review of linguistic behaviour by Jonathan Bennett. *Language, 53*(2), 417–424.
Kutz, C. (2000). Acting together. *Philosophy and Phenomenological Research, 61*(1), 1–31.
Ludwig, K. (2007). Collective intentional behavior from the standpoint of semantics. *Nous, 41*(3), 355–393.
Ludwig, K. (2013). Proxy agency in collective action. *Nous*. doi:10.1111/nous.12013.
Mele, A. R. (1992). Recent work on intentional action. *American Philosophical Quarterly, 29*(3), 199–217. doi:10.2307/20014416.
Recanati, F. (1986). On defining communicative intentions. *Mind & Language, 1*(3), 213–241. doi:10.1111/j.1468-0017.1986.tb00102.x.
Rubinstein, A. (1989). The electronic mail game: strategic behavior under "almost common knowledge". *The American Economic Review, 79*(3), 385–391. doi:10.2307/1806851.
Schiffer, S. R. (1972). *Meaning*. Oxford: Clarendon Press.
Searle, J. R. (1990). Collective intentions and actions. In P. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication* (pp. 401–415). Cambridge: MIT Press.
Sperber, D., & Wilson, D. (1996). *Relevance: Communication and cognition*. Oxford: Wiley.
Strawson, P. F. (1964). Intention and convention in speech acts. *Philosophical Review, 73*(4), 439–460.
Tuomela, R. (1991). We will do it: An analysis of group-intentions. *Philosophy and Phenomenological Research, 51*(2), 249–277.
van der Grijn, F. (1994). (Im)possibility of a coordinated attack *ILLC Publications: Preprints and Other Reports*: Universiteit van Amsterdam.