Explaining Speaker/Audience Coordination

Richard Warner*

Contents

"Pragmatics involves perception augmented by some species of 'ampliative' inference—induction, inference to the best explanation, Bayesian reasoning, or perhaps some special application of general principles special to communication, as conceived by Grice . . .—but in any case a sort of reasoning."[1] Pragmatics (at least a good deal of it) assumes that one's competence as a reasoner and user of natural language is sufficient to allow one to construct, in a range of significant cases, plausible accounts of how speakers and audiences reason.[2] The problem is that it

* Professor of Law, Chicago-Kent College of Law, Visiting Foreign Professor, University of Gdańsk, Poland.

[2] See, e. g., K.M. Jaszczolt, "Post-Gricean Pragmatics" in L. Cummings (ed.), The Routledge Pragmatics Encyclopedia. London: Routledge. 362-364.

appears that speakers and their audiences rarely engage in the reasoning pragmatics typically attributes to them. This sentence is an example. You understand it straightaway without any explicit reasoning, and indeed seemingly without any reasoning, explicit or otherwise. The same is true for me as the speaker: the words occurred to me as I wrote them—seemingly without any reasoning leading me to do so. The worry is a general—and old—one, as this 1918 critique of Jeremy Bentham's utilitarianism illustrates: "In the social sciences we are suffering from a curious mental derangement … the orthodox doctrines of economics, politics and law rest upon a tacit assumption that man's behavior is dominated by rational calculation … [even though] this is an assumption contrary to fact."[3] Are those who attribute reasoning to speakers and audiences "curious mental derangement" that prevents them from seeing the obvious fact that reasoning is rare?

I consider four replies. (1) Speakers and audiences do reason to the extent pragmatic explanations require; they just typically do not reason consciously. (2) The second reply concedes that speakers and audiences often do not reason even unconsciously in any relevant detail, but it insists that attributions of reasoning can nonetheless be, and often are, explanatory. (3) The third reply is a response to objections to the second. It identifies reasoning with information processing steps. (4) The fourth reply concedes defeat. It abandons the attempt to see speakers and audiences as reasoning. It explains communicative interactions simply by appeal

[3] Mitchell, Wesley Clair. 1918. "Bentham's Felicific Calculus." Political Science Quarterly 33(2): 161– 183.

information processing without any claim that the processing qualifies as reasoning. None of the four replies is adequate, so I suggest a fifth.

A review of speaker meaning is a preliminary.

Speaker Meaning

At an intersection, a driver you flashes her lights at you, making you wonder, "Why is she doing that?"[4] You answer, "She must intend me to believe that my lights are not on. She would not have that intention if they were on. So, they are not on." Your reasoning attributes to the other driver the following intention. She flashes her lights intending (1) that you believe that your lights are not on, and (2) that you recognize her intention (1), and (3) that that this recognition be part of your reason for believing that your lights are not on. Call such an intention an M-intention. For my purposes, the following simple definition of M-intentions suffices (using 'utter' and its cognates to cover both linguistic and non-linguistic items): for all propositions p,[5] a speaker S M-intends that p for an audience A by uttering U if and only if S utters U intending (1) that A believe that p; (2) that A recognize S's

[4] This section draws on Richard Warner, "Coordinating Meaning: Contractual Promises, Common Knowledge and Coordination in Speaker Meaning," in Alessandro Capone (ed.*), Further Advances In Pragmatics And Philosophy*, Springer (2018), and Richard Warner, "Indirect Reports and the Law," in Alessandro Capone (ed.), *Indirect Reports, Pragmatics, and the Law*, Springer, forthcoming.

[5] The notion of a proposition is fraught with difficulties. In addition toQuinean doubts about a notion that relies on synonymy relations for individuation, there are the problems Stephen Schiffer groups under "vagueness." Stephen Schiffer, "Gricean Semantics and Vague Speaker-Meaning," *Croatian Journal of Philosophy*, Vol. 16, 293, 2017. We could eliminate talk of propositions if we had a theory that explained how the brain combines lexical representations into every more complex structures in ways that underlie thought and language. The 'p' in 'S means p' could then be interpreted as ranging over those structures. Lacking such a theory, we have to make do with the promissory note drawn on such a theory—the concept of a proposition. The concept is a stopgap measure. As Schiffer notes, "What could be the point of trading in facts about meaning for facts about the content of beliefs if one ends up with nothing to say about the latter?" Stephen Schiffer, *Remnants of Meaning* (Cambridge, Mass.: The MIT Press, 1989), 2.

intention (1), and (3) that this recognition be part of A's reason for believing p. A speaker S means that p by producing X if and only if S M-intends that p by producing X.[6] The intersection example is a case of *asserting* that your lights are not on. Later work by Grice and others refines the analysis to meet a variety of counterexamples, and it extends the analysis to the full range of speech acts.

I set the refinements aside, and I also discuss only one speech act—assertion, where the intended response is belief. My focus is not on the definitional refinements, but on the underlying explanatory idea: namely, that speakers and the audiences coordinate utterance and response through complementary reasoning. I will refer to this as the "Gricean explanation." It is helpful to be more specific about the relevant type of coordination. To this end, let us say that speaker and an audience *M-coordinate* when and only when the speaker, by some utterance, M-intends that p for the audience, and the audience realizes that the speaker M-intends that p. The Gricean Explanation of M-coordination is that speakers and audiences engage complementary in reasoning about the speaker's intentions. Grice's account of speaker meaning gains much of its attractiveness and power from the Gricean explanation.

The Four Replies

I consider four replies to the objection that speakers and audiences do not reason in the relevant ways.

First: Unconscious Reasoning

[6] H. P. Grice, "Meaning," *Philosophical Review* 66 (1957): 337.

A first reply is that, in some cases at least, reasoning about M-intentions seems relatively unproblematic. Suppose Victoria says “Let’s get the kids something.”[7] With the children present, Victor responds by uttering ‘Okay, but I veto I-C-E C-R-E-A-M.’ Victoria realizes that he is spelling of ‘ice cream’ so that the children won’t realize that is what he is talking about. Why does she realize that? One plausible explanation is that Victor’s uttering ‘Okay, but I veto I-C-E C-R-E-A-M’ makes it common knowledge between them that ‘Okay, but I veto I-C-E C-R-E-A-M’ can be used to M-intend that one vetoes ice cream. Victoria then infers that Victor does indeed M-intend that he vetoes ice cream and is spelling of ‘ice cream’ to avoid the children realizing that he has that M-intention. The point on which this reply turns is that Victoria can reason in this way even if her beliefs are not *occurrent*.

An occurrent belief is a belief that is present to consciousness in—more or less—the way your belief that you are reading the following sentence will shortly be present to your consciousness: “I am now reading this sentence.” The example may suggest that having an occurrent belief requires that relevant *words* are present to consciousness. There is, however, no such requirement. Imagine you are eating bittersweet chocolate. Awareness of the bittersweet taste pervades your consciousness, and that awareness is the occurrent belief. The words ‘bittersweet’ do not run through your mind, nor any other relevant words.

One’s reasoning need not involve any occurrent beliefs. Suppose, for example, you are driving down a street when child darts out from between parked

[7] The example is from Stephen C. Levinson, *Pragmatics* (Cambridge University Press, 1993), 104.

cars directly into the path of your car. You believe that if you do not swerve, you will hit the child, and you believe that, all things considered, you should not hit the child. You conclude that you should swerve, and you do. Onlookers would readily attribute the non-occurrent beliefs to you to explain why you swerved—as indeed you would yourself. Suppose that, when you swerve, you sideswipe a car that was pulling out of a parking space. When the driver asks why you hit her, you would very likely attribute the beliefs to yourself to explain why. You might for example, say "I had to swerve to avoid the child." You most likely would not produce the stilted "I believed that if I did not swerve, I would hit the child, and I believed that, all things considered, I should not hit the child. So I concluded that I should swerve," but the plausible point of what you do say is to attribute those beliefs to yourself. So why not see speakers and audiences as reasoning in the ways the model requires, just not occurrently?

Indeterminacy is the problem as Grice notes in *Aspects of Reason*.[8] To use his example, suppose Jill reasons: "Jack broke his crown, but he is English; therefore, he will be brave." She thinks this intending the inference signaled by 'therefore' to be valid. Should we assume that there is some non-occurrent belief that mediates Jill's inference? There may be. Jill may have carefully considered the issue of English bravery and arrived at the conclusion that people of Jack's age and description are invariably brave. But suppose Jill has never given much thought to the bravery of the English. Nonetheless, as a result of education and acculturation, she is disposed to believe *something like* "All the English are brave," or "English are brave when they break their crowns," or "People of Jack's age and description are

[8] Paul Grice, *Aspects of Reason*, ed. Richard Warner (Oxford: Clarendon Press, 2005).

brave provided that they are also English," and so on. The disposition explains her inference even though she lacks any *determinate* belief linking being English and being brave.

Some may object that Jill's self-ascription of a belief may identify a unique suppressed premise. Imagine Jack asks her why she thinks his being English means he will be brave. She might reply "I was thinking that English are brave when they break their crowns." Does this show that Jill's inference was mediated by the specific, but non-occurrent belief that the English are brave when they break their crowns? That is implausible. To begin with Jill, may readily alter the belief she ascribes to herself. Imagine Jack responds, "*All* English? Really?", and Jill responds, "No, of course, not. Just English like you, or maybe just *probably* brave when they are like you, something like that." Further, Jill may *not* identify a unique option even initially. When Jack asks why she thinks why being English means he will be brave, she might reply, "I must have been thinking 'English are brave' or 'English like you are brave'—something like that." Jill's possible self-ascriptions of belief actually support the claim that she may lack a determinate belief linking being English and being brave.

Similar remarks hold for the swerve example. What exactly do you believe when you see the child dart out? There are several possibilities: "A child darted out, and so I believed I had to swerve to avoid hitting the child," "A child darted out, so I had no choice, I had to swerve to avoid injuring the child," "There was a child in the path of my car, so I turned to avoid the child," and so on. With a little ingenuity, one can produce many more alternatives. As with Jill, you are disposed to believe *something like* "A child darted out, and so I believed I had to swerve to

avoid hitting the child." The disposition explains your inference without any *determinate* belief in the background.

The explanation-by-disposition point causes problems for the ascriptions of reasoning in pragmatics. Consider a simple example such as Victoria's reasoning about Victor's utterance of 'Okay, but I veto I-C-E C-R-E-A-M.' Victor utters that and Victoria realizes that he M-intends that he vetoes ice-cream. She also realizes he is spelling out 'ice cream' so the children will not realize he is talking about ice cream. Here is one representation of how Victoria might (non-occurrently) reason: "Why is Victor uttering 'Okay, but I veto I-C-E C-R-E-A-M'? Most likely, he M-intends that he vetoes ice cream. But then why is he spelling out 'ice cream'? Most likely, he intends that the children will not realize he is mentioning ice cream." The problem is that there are many other equally plausible representations of Victoria's reasoning. Here is another. "Why is Victor uttering 'Okay, but I veto I-C-E C-R-E-A-M'? People spell words in the presence of children when they do not want the children to realize what they are talking about. Our children have already had way too much ice cream, so the best explanation of Victor's utterance is that he does not want the children to realize he is talking about ice cream while also M-intending that he vetoes ice cream." With a little ingenuity, one can produce a variety of options for Victoria's reasoning. Which one of the possibilities is Victoria's reasoning?

The *prima facie* plausible response is, "None of them." We do not need to identify a particular sequence of reasoning to explain Victoria's response. A sufficient explanation is that she is disposed to reason in *something like* various possible explicit reasoning sequences. The *disposition* that explains her response.

In general, explaining speaker-meaning interactions, appears only to require attributing appropriate dispositions to speakers and audiences. For speakers, the disposition is to infer that the audience will believe a certain proposition if the speaker utters a certain sentence. For audiences, the disposition is to infer from the fact that a speaker uttered a certain sentence that the audience should believe a certain proposition. M-coordination coordination is not a matter of complementary reasoning, but of complementary dispositions.

Second: Reason-Preserving Transitions

The second reply is a response to the "the disposition explains" objection to the first reply. The reply distinguishes two types of dispositions. To illustrate the first type, imagine Roger, a hapless undergraduate, is disposed to assert "Chickens run around after their heads are cut off" if someone asks him to produce an argument that the human soul is immortal. Unfortunately for Roger, that is precisely what his professor does ask him to do.[9] Roger blurts out, "Chickens run around after their heads are cut off," but, when asked how headless chicken behavior supports the immortality of the human soul, he says he has no idea. To illustrate the second type of disposition, turn to Jill's "He is English, therefore he will be brave." Like Roger, Jill is disposed to make a certain belief-transition—in her case, from "Jack is English" to "Jack will be brave." Unlike Roger, Jill's transition is not just a leap of thought but—we may suppose—an instance of a *reason-preserving transition*. A transition is *reason-preserving* if and only if, necessarily, if one has reasons for the initial set, then one does for the subsequent set as well.

[9] The example is Grice's in Grice, *Aspects of Reason*.

The second reply claims that people have the capacity to make reason-preserving transitions, and speakers and audiences typically exercise that capacity in coordinating in instances of speaker meaning. So, it is true that dispositions explain coordination between a speaker and an audience, but the activation of those dispositions are exercises of the capacity for reason-preserving transitions. They are, in this sense, instances of reasoning.

Grant, for sake of argument, that all this is true. The reply nonetheless fails to explain speaker-audience coordination. To see why, recall that the first reply appeals to a series of reasoning steps to explain why a particular utterance produces a particular response. The second reply avoids that appeal, but it puts nothing else in its explanatory place. Suppose, for example, that Roger utters 'Victoria brandished her clarinet like a tomahawk" and in response Ralph believes that she did. The second reply merely asserts that Ralph's belief is the result of a reason-preserving transition from Ralph's belief that Roger made the utterance. It does not explain why *that* utterance yields a transition to *that* belief. Suppose, by way of analogy, a young child asks to explain why the water wheel turns when the water strikes it. It would not do to answer, "The water strikes the wheel, and the wheel's reaction is to turn." The child can see that much. The child wants to know what connects the striking of the water with the turning of the wheel. One answer would be that the wheel can freely turn around its axle and the weight of the water on one side but not the other makes it do so. The second reply provides no such event-connecting explanation for speaker-audience coordination. That is a high price to pay—at least for anyone who finds that explanation attractive, and it is a prohibitively high price for pragmatics—at least for those committed to the

methodological stance that one can, in a wide range of case construct plausible accounts of how speakers and audiences reason.

Third: Reasoning as Information Processing Steps

The third reply defends the view that particular reasoning steps explain M-coordination through an appeal to the biological endowment that consists in the capacity to learn and use natural languages. That endowment enables people to generate a potentially infinite number of semantically interpreted syntactic structures that serve as the content of a variety of psychological attitudes, and it also enables them to map those structures to the phonetic representationsband map phonetic representations back on to the structures.[10] We do not yet have an adequate *theory* of such information processing, but the following assumptions are plausible. (a) Human information processing uses (at some points) structures *something like* the syntactic structures of natural languages, where elements in those structures have *something like* the semantic properties of words, phrases, and sentences of natural languages. (b) There are, however, a number of transformations of various types of encodings before we get to the level of speech and explicit thought. (c) At some point, one or more of these encodings underlie in some yet to be specified way speech and occurrent propositional attitudes.

The third reply seizes on (a) – (c) to suggest that speakers and audiences do reason in a particular sequence about M-intentions. The reply makes three claims: first, determinate sequences of information processing underlie communication'

[10] Chomsky has repeated emphasized this point. See, for example, Noam Chomsky, *Powers and Prospects: Reflections on Human Nature and the Social Order*, 1st PB Edition (Boston, MA: South End Press, 1999).

second, those sequences typically do not, at least in large part, consist of *occurrent* beliefs; third, the sequences are nonetheless instances of *reasoning*. An analogy motivates the reply. Suppose, for example, that five years ago you read a book on quantum mechanics, and believed that electrons are either right-handed or left-handed. You still believe that now even though you have not occurrently believed it for years. It is plausible—or at least grant for the sake of argument that it is plausible—to see your non-occurrent belief as realized in some sort of neural encoding. So why not see the reasoning pragmatic analyses attribute as similarly encoded?

The problem (the one relevant here) is that the reply offers a promissory note, not an explanation. It is a plausible speculation that human information processing uses states with *something like* the natural language syntactic structures some elements of which have *something like* natural language semantic properties. It is an additional but also plausible speculation that some of these encodings underlie *in some yet to be specified way* speech and occurrent propositional attitudes. But these speculations are not an explanatory theory. They are gestures toward an empirical theory we currently entirely lack. We don't have the explanation. We just have the idea that there may be an explanation along vaguely described lines. Relying on a speculative vision of a future empirical theory is inconsistent with pragmatics methodological assumption that one's competence as speaker is sufficient to allow one to construct, in a range of significant cases, plausible accounts of how speakers and audiences reason.

Fourth: Abandon Ship

The fourth option is to abandon the Gricean Explanation and embrace the idea that explaining M-coordination requires an empirical theory of information processing, a theory that we do not currently have. That option looms large unless we can find a way to salvage the Gricean Explanation. The fifth reply is my suggestion.

The Fifth Reply: Common Knowledge and Coordination

The fifth reply appeals to common knowledge of M-intentions to explain M-coordination. People have common knowledge that something is the case if and only if they know it, know they know it, know they know they know it, and so on. I adopt Stephen Schiffer's treatment of speaker meaning in *Meaning*, which makes common knowledge part of the definition. Thus (roughly, using our highly simplified definition of M-intending, not Schiffer's sophisticated one): Speaker S M-intends that p for an audience A by uttering U in circumstances C if and only if S intends by uttering U in C to make it the case that it is *common knowledge* between S and A that S utters U intending (1) that A believe that p; (2) that A recognize S's intention (1), and (3) that this recognition be part of A's reason for believing p.[11]

I make three claims. (1) People typically have common knowledge of M-intentions. (2) Common knowledge of M-intentions is sufficient for M-coordination without any detailed reasoning about another's beliefs. (3) But practical and theoretical attributions of reasoning nonetheless play an explanatory role.

[11] See Stephen Schiffer, *Meaning* (Oxford: Oxford University Press, 1973). Schiffer's interesting analysis is much more complicated as it addresses a wide range of possible counterexamples and handles all types of speech acts.

To show that people have common knowledge of M-intentions it is first necessary to show common knowledge in general arises.

Common Knowledge

An example from the game theorist Michal Chwe illustrates how common knowledge can arise. During a game in 1996,

> baseball fans at Cleveland's Jacobs Field [looked] up to see an airplane pulling a banner advertising anonymous HIV testing. Obviously the irony here is the airing of such a sensitive issue as AIDS publicly and even festively on a bright sunny day at the ballpark . . . [The underlying purpose is that] I would be more likely to get an HIV test if I knew that doing so was not unusual, but I wouldn't find this out through everyday conversation; at the ballpark, looking up at the plane, however, it is obvious to all that everyone is seeing the same thing.[12]

It is obvious to all—everyone paying sufficient attention—that "everyone is seeing the same thing." More fully, (1) everyone (everyone paying sufficient attention) sees and thereby comes to know the banner is flying over the stadium, and (2) everyone knows (1). (1) and (2) give rise to the infinite sequence of knowledge levels that constitute common knowledge.[13]

To see how, suppose Alice and Bob are sitting together and looking up at the sign.[14] The claim is that (1) and (2) generate the following infinite sequence. Using subscripts to keep track of knowledge levels:

---

[12] Michael Suk-Young Chwe, *Rational Ritual: Culture, Coordination, and Common Knowledge* (Princeton, NJ: Princeton University Press, 2013), 41.

[13] This material is adapted from Richard Warner, "Coordinating Meaning: Contractual Promises, Common Knowledge and Coordination in Speaker Meaning," in ALESSANDRO CAPONE (ED.), FURTHER ADVANCES IN PRAGMATICS AND PHILOSOPHY, Springer, forthcoming; Robert H. Sloan and Richard Warner, "The Self, the Stasi, and the NSA: Privacy, Knowledge, and Complicity in the Surveillance State," *Minnesota Journal of Law, Science, and Technology* 17 (2016): 347.

[14] Assume each knows the other has his or her open-eyed head pointed toward the sign. More generally, let G be the group of people paying sufficient attention. Suppose x and y know that they are in G (x knows x is in G and y is in G, and y knows y is in G and x is in

First level: Alice knows$_1$ that Bob sees the sign.
Bob knows$_1$ that Alice sees the sign.

Second level: Alice knows$_2$ Bob knows$_1$ that Alice sees the sign.
Bob knows$_2$ Alice knows$_1$ that Bob sees the sign.

Third level: Alice knows$_3$ Bob knows$_2$ Alice knows$_1$ Bob sees the sign.
Bob knows$_3$ Alice knows$_2$ Bob knows$_1$ Alice sees the sign.

.
.
.

How do Alice and Bob take this infinite series of steps?

It is not difficult to see how they could take the first step. Alice can figure out that Bob sees the sign by reasoning this way: "I see Bob looking up toward the sign. He has normal perceptual abilities, so I can conclude that he sees the sign." Bob can reason the same way about Alice. This would yield the first level: Alice knows$_1$ that Bob see the sign, and Bob knows$_1$ that Alice sees the sign. They can get to the second level by reasoning about each other's reasoning at the first level. Alice can reason: "At the first level, I started from the fact I saw Bob's looking at the sign, and I reached the conclusion that Bob sees the sign. Now Bob sees me looking at the sign, and he has normal perceptual and reasoning capacities, so he will have reasoned just as I did to *his* conclusion that I see the sign."[15] For Alice to realize that fact about Bob is for her to reach the second-level conclusion: "I know$_2$ Bob knows$_1$ I see the sign." Bob can reason in the same way to his second-level conclusion that he knows$_2$ Alice knows$_1$ he sees the sign. Thus: Alice knows$_2$ Bob knows$_1$ Alice sees the sign, and Bob knows$_2$ Alice knows$_1$ Bob sees the sign. Alice

---

G). Then the x and y can generate an "x and y version" of the infinite sequence below. The focus on Alice and Bob is just for convenience.

[15] I have suppressed complications about the assumptions about reasoning abilities. See Ronald Fagin et al., *Reasoning About Knowledge*, Mit Press Paperback ed. edition (Cambridge, Mass: A Bradford Book, 2004), 6–7.

and Bob can get to the rest of the levels the same way they get from the first level to the second: by reasoning about their reasoning at the level below. For any level *n*, Alice reasons about Bob's *n* – 1 level reasoning to reach the conclusion that Alice $knows_n$ Bob $knows_{n-1}$ . . . that Alice sees the sign, and Bob reasons in the same way to reach the conclusion that that Bob $knows_n$ Alice $knows_{n-1}$ . . . that Bob sees the sign.

Alice and Bob *could* reason this way—for a few levels of knowledge, that is. No one can in a finite amount of time complete the steps required to generate knowledge throughout the *infinite* number of levels. Even at the lowest levels, reasoning seems unlikely. Consider the second level. How likely is it that Alice actually reasons (even unconsciously) this way: Alice can reason: "At the first level, I started from the fact I saw Bob's looking at the sign, and I reached the conclusion that Bob sees the sign. Now Bob sees me looking at the sign, and he has normal perceptual and reasoning capacities, so he will have reasoned just as I did to *his* conclusion that I see the sign"? One could argue over the likelihood of reasoning at lower levels, but there is a far more attractive alternative. As Thomas et. al. observe in their groundbreaking study, *The Psychology of Common Knowledge*, "common knowledge has its own representation, it need not contain multiple levels of embedded knowledge; it could consist of a single mental symbol which means, 'We have common knowledge.'"[16]

[16] Kyle A. Thomas et al., "The Psychology of Coordination and Common Knowledge," *Journal of Personality and Social Psychology* 107, no. 4 (2014): 13. There are two other experimental studies. Kyle A. Thomas et al., "Recursive Mentalizing and Common Knowledge in the Bystander Effect," *Journal of Experimental Psychology. General* 145, no. 5 (May 2016): 621–29, https://doi.org/10.1037/xge0000153. Peter M. Krafft et al., "Modeling Human Ad Hoc Coordination," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16 (Phoenix, Arizona: AAAI Press, 2016), 3740–3746, http://dl.acm.org/citation.cfm?id=3016387.3016430.

It is in fact routine for a representation (a “mental symbol”) to encode a large number of steps. You know for example that, starting on the first step of a 100 step staircase, it takes 99 steps to get to the top. You don’t count from 1 to 100 to figure this out. You use this representation: “100 – 1 = 99.” It is also routine for a representation to encode the ability to take an infinite number of steps (setting aside the limit of time)—for example: “add 1 to 0 and to each result of adding 1.” So what representation could encode the infinite number of steps in common knowledge? A representation of a process as meeting these two conditions: (1) the process results in all members of a group G knowing X, and (2) the process ensures that all members of G know that all members of G undergo the process and as a result come to know X. As the “sign over the stadium” example shows, knowing (1) and (2) gives you the ability to generate an infinite series of levels of knowledge. Call a representation that meets these conditions a *common knowledge generator*.

The sign over the stadium is one example of a common knowledge generator. Making eye-contact is another. Imagine that Sally sees Roger, an old acquaintance. She stares at him hoping to remember his name before he turns his head and sees her. Unfortunately for Sally, Roger turns his head too soon, and Sally realizes that it would be pointless to pretend she did not see him. It is pointless because it is now common knowledge between them that they see each other.

Common Knowledge and Coordination

Common knowledge is not necessary for people to coordinate their behavior,[17] but it greatly facilitates their coordination. Indeed, as Kyle *et al*. note, "[a]ctors coordinate when they have evidence for common knowledge, and refrain from coordinating when they do not."[18] Common knowledge facilitates coordination by making the parties *transparent*. Everything is out in the open, so there is no possibility of misunderstanding, misinterpretation, doubt, or deception at any knowledge level.[19] To illustrate the point, suppose that Victor and Victoria are discussing by cell phone whether to meet at the opera later in the evening or whether each will stay home alone.[20] Before they decide, their batteries run out, and they have no other way to communicate. They both have, and know they both have, the following preferences in the following order. (1) Attend the opera together. There is a benefit they hope to achieve thereby: namely, the pleasure of each other's company. (2) Stay home alone when the other does too. (3) Stay home alone when the other goes to the opera. (4) Go to the opera when the other does not. Given these preferences, Victor will go to the opera if he thinks Victoria will, and that he will stay home if he thinks she will stay home too. Vice versa for Victoria. Relevant common knowledge makes it easy for Victor and Victoria to coordinate their actions by both going to the opera or both staying home. Suppose, for example, they have a prior agreement that, when communication about going to the opera breaks down before a decision, they both go to the opera (or both stay

[17] For a fuller discussion, see Sloan and Warner, "The Self, the Stasi, and the NSA: Privacy, Knowledge, and Complicity in the Surveillance State."
[18] Thomas et al., "The Psychology of Coordination and Common Knowledge," 671.
[19] See Stephen Schiffer, *Meaning* (Oxford: Oxford University Press, 1973).
[20] Sloan and Warner, "The Self, the Stasi, and the NSA: Privacy, Knowledge, and Complicity in the Surveillance State."

home). That agreement generates the common knowledge that each will go (or stay home), and each acts accordingly.

Common knowledge can explain M-coordination in communication contexts as long as people can have common knowledge of M-intentions in those contexts. They can if there are relevant common knowledge generators in those contexts. Schiffer's definition of speaker meaning assumes the existence of such generators. According to Schiffer:

> S meant something by (or in) uttering x iff S uttered x intending thereby to realize a certain state of affairs E which is (intended by S to be) such that the obtainment of E is sufficient for S and a certain audience A mutually knowing* (or believing*) that E obtains and that E is conclusive (very good or good) evidence that S uttered x intending
>
> (1) to produce a certain response r in A;
>
> (2) A's recognition of S's intention (1) to function as at least part of A's reason for A's response r;
>
> (3) to realize E.[21]

E is a common knowledge generator. Do such generators exist? That is an empirical question about whether people have representations of a process as one that meets these conditions: (1) the process results in all members of the relevant group G knowing that a speaker make a specific utterance in a certain context has a relevant M-intention, and (2) the process ensures that all members of G know that all members of G undergo the process and as a result come to know that a speaker who makes that utterance in that context has the M-intention. The obvious candidate for the process is the acquisition of a natural language, with the relevant group being competent speakers of that language. It is plausible that utterances in particular circumstances of sentences (words, phrases) of the language satisfy relevant instances of (1) and (2).

[21] Schiffer, *Meaning*. 'Mutual knowledge*' is Schiffer's term for common knowledge.

A Role for Reasoning

So far, the fifth alternative looks more like an objection to the Gricean Explanation rather than its salvation. Relevant common knowledge explains M-coordination, and that common knowledge can exist without reasoning about others' psychological attitudes. It just requires the existence of the relevant common knowledge generator. So what explanatory role is left for providing accounts of how speakers and audiences could reason their way to M-coordination?

I submit there is. Providing such an account is one way to show that a relevant common knowledge generator exists. Consider again Victor's uttering 'Okay, but I veto I-C-E C-R-E-A-M.' His utterance makes it common knowledge between them that 'Okay, but I veto I-C-E C-R-E-A-M' can be used to M-intend that Victor agrees with getting the kids something but that he vetoes ice cream. The relevant common knowledge generator is the utterance of the sentence 'Okay, but I veto I-C-E C-R-E-A-M.' But Victor and Victoria's common knowledge does not stop there. Victor's utterance also makes it common knowledge (conversationally implicates, if you like) that he is spelling 'ice cream' so the children will not realize that he is mentioning it. What is the relevant common knowledge generator?

One way to identify a generator is to produce a sample of the reasoning it would generate. It is easy to do so in the ice cream example. Before his utterance, Victor could reason: "People spell things in front of small children when they don't want the children to know what they are talking about. So if I spell out 'ice cream,' Victoria will realize that I intend that the children should not realize I am mentioning ice cream." Victoria can easily replicate Victor's reasoning, and for any

knowledge level n, each can replicate the other's reasoning at level *n* – 1. This reasoning depends on background about adults' linguistic practices and on contextual knowledge (the presence of the children and so on). Those features most likely comprise the common knowledge generator. In other cases, constructing replicable reasoning my highlight significant syntactic and semantic features.

The Letter, Not the Spirit?

My "salvation" of the Gricean Explanation preserves an explanatory role for step by step representations of reasoning but abandons the idea that particular reasoning steps connect utterance and response. "I am conscious that my version, however close to the letter, is very far from the spirit of the original theory; but to defend the spirit as well as the letter would be beyond my powers."[22]

[22] Grice H. P. 'The Causal Theory of Perception ," Proceedings of the Aristotelian Society, Supp. vol. nqv. 1961. P. 152