

Stalnaker on Assertion

Intention & Acts of Meaning Seminar, Week 10

Daniel W. Harris

April 15, 2016

1 Why Are We Studying Stalnaker?

- His model of conversation is probably the most influential one around these days.
- Highly ranked philosophy departments are densely packed with his students (and their students). They (and many others) take central components of his theory of conversation for granted. So it's super important to be able to engage with Stalnakerian assumptions in your work.
- He sometimes suggests that his theory is an implementation of Griceanism, or that it is inspired by Grice's ideas⁴. The details of how Stalnaker relates to Grice are pretty interesting.

2 Important Background Assumptions

Semantic Content: The semantic value of a sentence relative to a context (or an utterance of a sentence) is a proposition.

- What facts about a context determine context-sensitive expressions' semantic contents? Stalnaker doesn't say much about this. But he assumes that this role can be played by a possible world. (He later says we need centered worlds to play this role.)

Possible-Worlds Propositions A proposition is identical to the set of worlds in which it is true. The proposition that Grice was right is the set of worlds in which Grice was right.¹

- This idea has a lot of problems: it collapses all necessary truths and all necessary falsehoods. It raises the problem of logical omniscience, it gives rise to loads of Frege puzzles, etc.
- But it also has a lot of practical benefits: it makes it easy to do mathematical operations on propositional contents, and so it is usually assumed in compositional semantics, formal epistemology, modal logic, etc.
- Some, including a lot of Stalnaker's students, will tell you that possible-worlds propositions are just an extremely useful idealization. We are looking at an intentionally distorted notion of content because doing so allows us to do useful formal tricks. We'll figure out how to build the extra fineness-of-grain back in later.
- Stalnaker can't say this. In *Inquiry* (1984), he defends a version of functionalism about the contents of mental states and argues that if this version of functionalism is right, then the contents of mental states can't be any more fine-grained than sets of worlds.
- Stalnaker's own attempts to solve the problems with possible-worlds propositions are pretty unsatisfying (we'll look at one of them below), but luckily he's got sympathetic colleagues at MIT. For some high-tech attempts to solve the problems, see Elga and Rayo (2015); Yablo (2014). Recently, Stalnaker has seemed to endorse Elga and Rayo's new solution.

values. For example, the proposition that Grice was right could be modeled as $[\lambda w_{st} . \text{Grice was right at } w]$ —a function that maps a possible world to 1 (truth) iff Grice was right at that world. This is strictly equivalent to the set-of-worlds idea, and although it seems more complicated at first, it is convenient to work with intensions when doing semantics, and so you'll see them used there.

¹This is equivalent to the idea that a proposition is an intension—a function that maps worlds to truth

3 Stalnaker’s Model of Conversation, Briefly

Stalnaker’s model of conversation is built around an idealized model of conversation:

- A conversation involves a determinate number of *participants* (a.k.a. *interlocutors*).
- The *context* is a body of information that the participants take for granted in for the purposes of the conversation. Stalnaker models this body of information in two equivalent ways:
 - The *common ground* is a set of propositions. Treating each proposition as a set of possible worlds (see below), a common ground can be modeled as a set of sets of worlds.
 - The *context set* is a set of possible worlds—the set of worlds that are in the intersection of all the propositions in the common ground. Intuitively, this is the set of worlds compatible with the information in the common ground.
- The goal of a conversation is to add information to the common ground. (Conversation of this kind is sometimes called ‘inquiry’ or ‘joint inquiry’.)
- By adding information to the common ground, we also reduce the number of worlds in the context set. So the aim of inquiry is to narrow down the context set to as few worlds as possible. In normal circumstances, the goal is to zero in on the actual world—to find out the truth.
- One way to add information to the common ground (thereby culling the context set) is to *assert* a proposition.
- Stalnaker says that the *characteristic effect* of an assertion is to add information to the common ground.
- In some places he suggests that aiming at this effect is what *makes* a speech act an assertion (more below).

Discussion Question: What is idealized about this model?

4 Theoretical Roles of Context

Stalnaker and others have posited common ground to play many roles. I’ll focus on three such roles: in metasemantics, the epistemology of interpretation, and in speech-act theory.

4.1 Metasemantics

The content of a context-sensitive expression is determined in part by the context in which it is uttered. Stalnaker doesn’t spell out how this works, but takes the context of utterance to be some feature of the possible world in which an utterance is produced.

Stalnaker thus says that there are two ways in which “the facts enter into the determination of a truth value of what is expressed in an utterance” (Stalnaker, 1999, 80)—as the context relative to which a proposition is expressed, and as the circumstances relative to which the truth of that proposition is evaluated. He illustrates this idea by discussing a disagreement he once had with a pair Irish-Americans:

Let me give a simple example. I said *You are a fool* to O’Leary. O’Leary is a fool, so what I said is true, although O’Leary does not think so. Now Daniels, who is no fool, and who knows it, was standing near by, and he thought I was talking to him. So both O’Leary and Daniels thought I said something false: O’Leary understood what I said, but disagrees with me about the facts; Daniels, on the other hand, agrees with me about the fact (he knows O’Leary is a fool), but misunderstood what I said. Just to fill out the example, let me add that O’Leary believes falsely that Daniels is a fool. Now compare the possible worlds *i*, *j*, and *k*. *i* is the world as it is, the world we are in; *j* is the world that O’Leary thinks we are in; and *k* is the world Daniels thinks we are in. If we ignore possible worlds other than *i*, *j* and *k*, we can use matrix A to represent the proposition I actually expressed.

(A)

<i>i</i>	<i>j</i>	<i>k</i>
T	F	T

But the following TWO-DIMENSIONAL matrix also represents the second way that the truth-value of my utterance is a function of the facts:

(B)

.	<i>i</i>	<i>j</i>	<i>k</i>
<i>i</i>	T	F	T
<i>j</i>	T	F	T
<i>k</i>	F	T	F

The vertical axis represents possible worlds in their role as context—as what determines what is said. The horizontal axis represents possible worlds in their role as the arguments of the functions which are the propositions expressed. Thus the different horizontal lines represent WHAT IS SAID in the utterance in various different possible contexts. (Stalnaker, 1999, 80–1)

4.2 Interpretation

It is widely held that the common ground is the body of information that interpreters do or should rely on when interpreting speakers, and that speakers should assume hearers will take into account when interpreting them.

Here's a particularly strong version of this thesis:

...for a listener to understand a speaker's meaning, he can confine himself to a certain limited domain of information, namely, the speaker's and his listener's common ground, that part of the speaker's and listener's knowledge, beliefs and assumptions that are shared. (Clark and Carlton, 1992, 61)

Stalnaker seems to assume something similar (see, e.g., the last paragraph of §2.2 of *Context* 2014, 42) For objections to this view, see (Sperber and Wilson, 1995, §1.3).

Stalnaker's theory of diagonalization gives context another privileged role to play in interpretation. (See below.)

4.3 Speech Acts

Context plays a third kind of role, for Stalnaker: it is the target of speech acts. In some sense, that is, it is the *essential effect* of a speech act to change the context.

His only example is assertion. The essential effect of asserting p is to add p to the common ground. (Equivalently, it is to intersect p with the context set.) As we'll see next week, others have proposed analogous theories of other speech acts.²

How is the performance of a speech act defined in terms of its essential effect? Some options:

- A simple answer: assertions are just those speech acts that *actually* add their content to the common ground.
 - But this won't work, because some assertions are rejected or misunderstood.
- In 'Assertion' (1978), Stalnaker therefore initially suggests that asserting p is a speech act that either (a) succeeds in adding p to the common ground, or (b) is rejected, but would have added p to the common ground if it hadn't been rejected.
 - But Stalnaker makes it clear that he does not "propose this as a *definition* of assertion, but only as a claim about one effect which assertions have, and are intended to have—an effect that should be a component, or a consequence, of an adequate definition" (1999, 87).

²Sometimes people talk about "characteristic effects" instead of "essential effects".

- One problem: the account does not provide a necessary condition on assertion, since "there may be various indirect, even nonlinguistic, means of accomplishing the same effect which I would not want to call assertions" (1999, 87).
- Another problem: this account "makes reference to another speech act—the rejection of an assertion, which presumably cannot be explained independently of assertion" (1999, 87).
- Also, it seems pretty gerrymandered.

- In more recent work, Stalnaker tends to define asserting p as *proposing* to add p to the common ground:

...assertion should be understood as a proposal to change the context by adding the content to the information presupposed. This is an account of the force of an assertion.... (Stalnaker, 1999, 10–11)

...the illocutionary force of an assertion (on the kind of account I have promoted) is explained as a proposal to change the common ground in a certain way. (Stalnaker, 2014, 39)

But what is *proposing*? It sounds like yet another speech act! So this is open to the same objection as before.

- Another option: to assert p is to intend (or, better, communicatively intend) to add p to the common ground. A prominent defense of this is by Thomason (1990):

To mean p is to intentionally reveal an intention to [add p to the common ground] through the hearer's recognition of the status of an intention or plan of the speaker's. (Thomason, 1990, 345)

Stalnaker gives the following objection to making assertion a matter of what one intends:

My suggestion about the essential effect of assertion does not imply that speakers *intend* to succeed in getting the addressee accept the content of the assertion, or that they believe they will, or even might succeed. A person may make an assertion knowing it will be vetoed, a labor negotiator may make a proposal knowing it will be met by a counterproposal, or a poker player may place a bet knowing it will cause all the other players to fold. Such actions will not be pointless, since they all have secondary effects, and there is no reason why achieving the secondary effects cannot be the primary intention of the agent

performing the action. The essential effects will still be relevant even when it is a foregone conclusion that the assertion, legislative act, proposal, or bet will be rejected, since one generally explains why the action has the secondary effects it has partly in terms of the fact that it would have had certain essential effects had it not been rejected. (Stalnaker, 1999, 87)

A tempting reply: cases in which “secondary effects” are the real point, and in which the characteristic effect is not actually intended, are not genuine cases of assertion, but only cases of *making as if* to assert in order to (indirectly) do something else.

5 The Nature of Common Ground

What makes it the case that the common ground of a conversation contains a given proposition (at a given time)? Stalnaker’s theory about this has evolved.

5.1 Stalnaker’s 1978 Theory of Common Ground

Stalnaker initially defines common ground in terms of presupposition:

Roughly speaking, the presuppositions of a speaker are the propositions whose truth he takes for granted as part of the background of the conversation. A proposition is presupposed if the speaker is disposed to act as if he assumes or believes that the proposition is true, and as if he assumes or believes that his audience assumes or believes that it is true as well. Presuppositions are what is taken by the speaker to be COMMON GROUND of the participants in the conversation, what is treated as their *common knowledge* or *mutual knowledge*. (Stalnaker, 1999, 84)

A proposition is in the common ground of a conversation, at a given moment, if all of the participants in the conversation presuppose it, at that moment. Some points worth noting:

- A context is *non-defective* if every participant in the conversation presupposes the same propositions.
- A context is defective if the participants’ presuppositions don’t line up.
- All real-world contexts are probably defective to some extent, but they are often defective in ways that don’t matter. They are “close enough”.
- In some cases, a context may be defective in ways that leads to miscommunication.

5.2 The Updated Theory (1984–2014)

In *Inquiry*, Stalnaker broadens his theory of the attitudes underlying common ground. Instead of just presupposition, he now talks about *acceptance*. Here’s how he explains the concept:

Acceptance, as I shall use the term, is a broader concept than belief; it is a generic propositional attitude concept with such notions as presupposing, presuming, postulating, positing, assuming and supposing as well as believing falling under it. Acceptance is a technical term: claims I make about acceptance are not intended as part of an analysis of a term from common usage. But I do want to claim that this technical term picks out a natural class of propositional attitudes about which one can usefully generalize. Belief is obviously the most fundamental acceptance concept, but various methodological postures that one may take toward a proposition in the course of an inquiry or conversation are sufficiently like belief in some respects to justify treating them together with it.

To accept a proposition is to treat it as a true proposition in one way or another—to ignore, for the moment at least, the possibility that it is false. One may do this for different reasons, more or less tentatively, more or less self-consciously, with more or less justification, and with more or less feeling of commitment. (Stalnaker, 1984, 79–80).

Worth noting:

- Unlike belief, agents don’t accept things *simpliciter*, but always “for some purpose or other” (Stalnaker, 2014, 39).
- This is because “acceptance may be the product of methodological decision rather than subjective commitment. One may accept something for the sake of argument, although one cannot believe things for this reason” (Stalnaker, 1984, 81).
- A consequence is that “a person may accept something in one context, while rejecting it or suspending judgment in another. There need be no conflict that must be resolved when the difference is noticed, and he need not change his mind when he moves from one context to the other” (Stalnaker, 1984, 80–81).
- When we talk of what the participants in a conversation mutually accept, therefore, this must be interpreted as what they mutually accept *for the purposes of the conversation*.

The common ground, according to Stalnaker (2014), is the set of propositions that the participants in a conversation mutually accept, for the purposes of the conversation. *Mutual* acceptance is defined on the model of the traditional, iterated definition of mutual knowledge, as follows:

MUTUAL ACCEPTANCE

For any agents A and B and any proposition p , A and B mutually accept p if and only if:

- (i) a. A accepts p ;
b. B accepts p ;
- (ii) a. A accepts that B accepts p ;
b. B accepts that A accepts p ;
- (iii) a. A accepts that B accepts that A accepts p ;
b. B accepts that A accepts that B accepts p ;

And so on, *ad infinitum*.

5.3 Context and As-If Attitudes

A common feature of both early and late Stalnaker is that context is constituted by *as-if* attitudes: like accepting p , presupposing p is neither necessary nor sufficient for knowing or believing p :

The propositions presupposed in the intended sense need not really be common or mutual knowledge; the speaker need not even believe them. He may presuppose any proposition that he finds it convenient to assume for the purpose of the conversation, provided that he is prepared to assume that his audience will assume it along with him. (Stalnaker, 1999, 84)

Why does context have this as-if quality? Well, imagine the following conversation:

SPIES

Eve and James are English spies and Felix is an American spy. They are all operating undercover: Eve is posing as an arms dealer, James as her assistant, and Felix as an interested buyer. Eve and James both know (and each knows that the other knows, etc.) that they are both spies. Felix knows that he (Felix) is a spy. Eve and James also know that Felix is a spy, and Felix knows that

they are spies. But neither Eve nor James knows that Felix knows that they are spies, and Felix likewise doesn't know that Eve and James know that Felix is a spy. In this context, we should expect that everyone will proceed as if their cover hasn't been blown. In particular, although none of our heroes knows or believes (mutually or otherwise) the proposition that James is Eve's assistant, we can still think of them as mutually accepting this proposition, since they will continue to treat it as true, for the time being, and for the purposes of this conversation. Suppose that in this context, Felix utters the following sentence:

1. Your assistant makes excellent cocktails.

How should we expect the discourse context (i.e., the common ground) to do its various purported jobs in this case?

6 Diagonalization

Suppose we're engaging in some inquiry about the presidential election. We're collectively wondering about two things: (a) whether Trump or Cruz would win the primary, and (b) whether a Republican or a Democrat will win the presidential nomination. So the following four worlds are in the context set:

- TR The world in which Trump wins the primary and a Republican (namely, Trump) wins the general election.
- TD The world in which Trump wins the primary but a Democrat wins the general election.
- CR The world in which Cruz wins the primary and a Republican (namely, Cruz) wins the general election.
- CD The world in which Cruz wins the primary but a Democrat wins the general election.

None of us knows what to think, so we go and consult Nate Silver. Lately, Nate has been using the name 'Republican Rodney' as a name for the eventual winner of the 2016 Republican primary.³ When we consult Nate, he performs an assertion with the following utterance:

³Specifically, suppose that 'Republican Rodney' is a rigid designator for whomever the Republican nominee will turn out to be, much like 'Julius' is a rigid designator for whoever invented the zip.

(X) Republican Rodney will lose the general election.

Which worlds should we eliminate from the context set as a result of Silver’s oracular assertion? Which worlds are no longer compatible with what we’re mutually accepting?

- Well, given that we don’t (yet) know who Republican Rodney is, we can’t eliminate TD or CD.
- But we can eliminate all worlds in which a Republican wins the general election—namely, TR and CR. Right?
- This means that the proposition that Silver asserted was {TD, CD}, or, more naturally, not-{TR, CR}. This is the proposition that whoever ‘Republican Rodney’ turns out to refer to, that guy will lose the general election.

Problem: that proposition cannot be the semantic content of the sentence Silver uttered, no matter what context we’re in! After all: Republican Rodney is a rigid designator. So the propositional concept of (X) is given (let us suppose) as (Y):

	TR	TD	CR	CD
TR	F	T	F	F
TD	F	T	F	F
CR	F	F	F	T
CD	F	F	F	T

The proposition {TD, CD} is not expressed at any context (i.e., on any row). So how can Silver have asserted this proposition by uttering (X)?

- Stalnaker noticed that the proposition intuitively asserted in this case is the *diagonal* proposition—the one we get by moving from the upper left cell of (Y) diagonally down to the lower right cell.
- In general, the diagonal proposition for a sentence S (written †S) will have the following truth conditions: for each world w in the context set, the proposition that S expresses relative to w is true at w . In effect, diagonalizing turns all of the rigid designators in a sentence into metalinguistic descriptions.
- Stalnaker’s proposal is that we diagonalize as a kind of “repair strategy”—a pragmatic fallback when there is something semantically wrong with the sentence that’s been uttered, relative to the context in which it’s been uttered.

- So what’s wrong with (X) in this context? Intuitively, it’s that we don’t know who ‘Republican Rodney’ refers to, and so we can’t figure out what the literal semantic content is. In situations like that, we substitute the diagonal proposition for the semantic content.

Stalnaker proposes the three principles, which he compares to Gricean maxims, which, if violated (flouted?), trigger diagonalization. He says of these principles:

...Given the framework of propositions, presupposition, and assertion, the principles are all pretty obvious. They are not intended as empirical generalizations about how particular languages or idiosyncratic social practices work. Rather, they are proposed as principles that can be defended as essential conditions of rational communication, as principles to which any rational agent would conform if he were engaged in a practice that fits the kind of very abstract and schematic sketch of communication I have given. (Stalnaker, 1999, 88)

Here are the principles, together with some intuitive motivation for them.

(S1) A proposition asserted is always true in some but not all of the possible worlds in the context set.

- Asserting something that’s already in the common ground would be redundant, so don’t do that.
- Asserting something that is incompatible with the common ground would be impossible for others to process, so don’t do that!

(S2) Any assertive utterance should express a proposition, relative to each possible world in the context set, and that proposition should have a truth value in each possible world in the context set.

- Don’t use sentences that presuppose things that you’re not presupposing.
- In particular, don’t use a referring expression if it’s not common ground that it actually refers to anything. Otherwise people won’t know what to make of what you’re saying.
- For example, don’t use ‘he’ at the start of a conversation without giving your hearers some indication of who you’re referring to.

(S3) The same proposition is expressed relative to each possible world in the context set.

- You shouldn't use a sentence if the context doesn't give your hearers enough information to understand it.
- For example: don't utter 'you' if it's unclear whether you're talking to O'Leary or Daniels, because nobody will be able to figure out what you've said.

If someone does produce an utterance that seems to violate one of S₁–S₃, and the context is non-defective, and they're being cooperative, then what they've really done is flouted the rule in question in order to communicate the diagonal proposition. That's what Nate Silver did in uttering (X).

7 Stalnaker and Grice

In 'Assertion', Stalnaker makes several shout-outs to Grice. E.g., in the acknowledgements:

The influence of Paul Grice's ideas about logic and conversation will also be evident. (Stalnaker, 1999, 78, fn)

Later, Stalnaker more explicitly credits Grice with the general strategy embodied in the principles that trigger diagonalization:

The influence of Grice's theory of conversation should be clear from my discussion of the application of these principles. (Stalnaker, 1999, 88, fn)

In §2.2 of *Context*, Stalnaker makes it sound as if his theory is meant to be understood in terms of Grice's broader framework.

In a footnote in 'Assertion', Stalnaker suggests that he got the idea of common ground from Grice (as well as from others):

Paul Grice spoke, in the William James Lectures, of propositions having *common ground status* in a conversation. (Stalnaker, 1999, 84, fn)

And at the end of §2.2 of *Context*, Stalnaker seems to say that something like his picture of common ground follows from some of Grice's views.

The upshot of this for our purposes is that common interest and common knowledge are necessary for the possibility of communication. Only against a relatively rich background of common belief is it possible to get people to recognize the very specific intentions that must be recognized for successful

acts of meaning, and only where there are mutually recognized common interests will the recognition of the intentions be effective in changing beliefs. (Stalnaker, 2014, 42)

This is a strong claim, and it's hard to see what, in the preceding pages, supports the details of it!

7.1 Grice on "Common-Ground Status"

Grice mentions "common ground" in his work on indicative conditionals. Roughly, he takes the natural-language indicative conditional to be material implication ('if p then q ' is semantically equivalent to $\neg p \vee q$).

But, he argues that it is normally felicitous to utter 'if p then q ' only if neither p nor $\neg p$ has "common ground status" in the conversation (Grice, 1989, 65–67, 84). Here's what he says:

- (5) [...] the negation of the antecedent of a conditional is to be thought of as common ground, as a possibility though not necessarily as a fact, between those who debate the pros and cons of a particular conditional.
- (6) The attribution of such common-ground status might be thought of either as something which has to be "read in" as a non-conventional implicatum, or as an element in the lexical meaning of certain conditionals, such as those the expression of which involves words or phrases such as "supposing," "suppose that," "if ... then," and perhaps even "if" itself.
- (7) It is further open to question, should it be necessary to attribute common-ground status, just what constitutes being common ground. There are two possibilities.
- (8) One is that it lies in an understanding by speakers that they are debating the nature of the acceptable alternative to that which is taken as being common ground as a possibility. ("If not so-and-so, then what?")
- (9) The other possibility would be one which involves the specification of a syntactical rule which would dictate the allocation of scope within a rewritten formulation of the original conditional.
- (10) If the second version is preferred, it will be difficult to avoid the supposition that it is a feature of the lexical meaning of the conditional, even though possibly a feature which rests upon prelexical considerations, and which introduces no new concepts. (Grice, 1989, 84)

Grice discusses “common-ground status” a bit elsewhere. In ‘Presupposition and Implicature’ (dated 1970/1977), he glosses the possession of common-ground status as “to be treated as noncontroversial” (Grice, 1989, 276) and “not to be questioned” (Grice, 1989, 278). In ‘Meaning Revisited’ (dated 1976/1980), he casually drops this in:

Here I am interested not so much in the existence of that distinction [between natural and non-natural meaning], which has now, I think, become pretty boringly common ground (or mutual knowledge), but rather in the relationship between the two notions, the connections rather than the dissimilarities between them. (Grice, 1989, 284)

Safe to say: Grice didn’t have a theory of common ground, and the notion did not play a major role in his theory of conversation!

7.2 Common-Ground-Like Notions in Schiffer

Schiffer argues that *m*-intentions must be spelled out in terms of mutual knowledge, which is one model for Stalnaker’s theory of presupposition and mutual acceptance.⁴

Schiffer makes this change in response to a kind of counterexample to Grice’s definition of speaker meaning raised by Strawson and then Schiffer, and that, according to Schiffer, leads to an infinite regress.

First, let’s remember Grice’s three-clause definition:

“*U* meant something by uttering *x*” is true iff, for some audience *A*, *U* uttered *x* intending

- (1) *A* to produce a particular response *r*
- (2) *A* to think (recognize) that *U* intends (1)
- (3) *A* to fulfill (1) on the basis of his fulfillment of (2).

Here’s Neale’s description of the regress problem:

The following type of example is used by Strawson and Schiffer to demonstrate that clauses (1), (2), and (3) do not specify a rich enough intention (or batch of intentions). Suppose *A*, a friend of mine, is about to buy a house. I think the house is rat-infested, but I don’t want to mention this outright to *A* so I let

rats loose in the house knowing that *A* is watching me. I know that *A* does not know that I know that he is watching me do this. I know *A* will not take the presence of my rats to be natural evidence that the house is rat-infested; but I do know, indeed I intend, that *A* will take my letting rats loose in the house as grounds for thinking that I intend to induce in him the belief that the house is rat-infested. Conditions (1)-(3) of (II) above are fulfilled. But surely it is not correct to say that by letting rats loose in the house I mean that the house is rat-infested.

The problem is that in this example my intentions are not, as Strawson puts it, wholly overt. One possible remedy involves adding a fourth clause:

- (4) *A* to recognize that *U* intends (2).

But as Strawson and Schiffer point out, with enough ingenuity the same sort of counterexample can still be generated, and then we need a fifth clause, then a sixth, and so on. (Neale, 1992, 549–50)

Schiffer’s solution was to reframe the definition of speaker meaning in terms of mutual knowledge. He has given many different versions (see, e.g., 1972, 50–8). Here’s a comparatively simple one:

S means* that *p* in uttering *x* only if, for some person *A* and relation *R*, *S* utters *x* intending it to be mutual knowledge* between *S* and *A* that *xRp* and, on the basis of this, that *S* uttered *x* with the intention of activating in *A* the belief that *p*. (Schiffer, 1982, 121)⁵

Stalnaker explicitly cites Schiffer, along with Grice (and Lewis), as one of the sources of the idea of common ground (Stalnaker, 1999, 84, fn). But there are some big differences between Schiffer’s use and Stalnaker’s:

- Stalnaker defines assertion as the attempt to change the common ground.
- Even after introducing mutual knowledge, Schiffer defines speech acts as attempts to change the mind (or actions) of some particular addressee.

⁵ Aside from mutual knowledge, two other features of this definition worth explaining: (i) Schiffer adds the asterisk to ‘means*’ as a reminder that he is giving a stipulative definition. (ii) Schiffer argues that, in order for *S* to mean something, there has to be some feature of one’s utterance that one intends *A* to use to recognize what one means. (Thus the role of *R* in his definition. This is originally discussed by Grice in ‘Utterer’s Meaning and Intentions’ as a response to certain kinds of counterexamples (Grice, 1989, 99–105).

⁴ Schiffer actually has a funky definition of mutual knowledge, which is something more like mutual justification. He therefore sometimes calls it ‘mutual knowledge*’ instead (Schiffer, 1972, §II.2).

- For Schiffer, mutual knowledge doesn't figure in the m-intended effect, only in what it takes to get that effect recognized.
- In other words: common ground does not play a role in Schiffer's replacement for Grice's clause (1); it only plays a role in Schiffer's replacement for Grice's clause (2).

This is followed by Grice's attempt to defuse Schiffer's worries about his definition of m-intentions leading to a regress. In the end, he suggests a solution different to Schiffer's.

7.3 Concepts like Mutual Knowledge

Most later Griceans after Schiffer put concepts like mutual knowledge somewhere in their theories of speaker meaning, linguistic convention, and so on. Bach & Harnish go for a weaker notion of mutual belief, for example:

A and B mutually believe that *p* if and only if each believes (1) that *p*, (2) that the other believes that *p*, and (3) that the other believes that the first believes that *p*. (1979, 285, n.4)

Even Sperber & Wilson, who are extremely critical of the idea of mutual knowledge, adopt the closely related notion, "mutual manifestness", and give it a role to play in their theory (1995). On their view, a proposition is mutually manifest to A and B iff it is perceptible or inferable from both present "cognitive context", and it is also perceptible or inferable that it is thus perceptible or inferable to the other. Other seminal discussion of mutual knowledge, common ground, and related notions include Clark and Carlton (1992); Clark and Marshall (1992); Lewis (1975)

But a problem: Harvey Lederman has recently come up with some persuasive arguments that nobody every has mutual knowledge, mutual belief, etc. (Lederman, 2015).

7.4 Grice on Mutual Knowledge

What did Grice think of mutual knowledge and related notions?⁶ Grice discusses the issue in 'Meaning Revisited'. Here's a telling passage:

...my general strategy was to look for the kind of regresses which Schiffer and others have claimed to detect concealed beneath the glossy surface of my writings on meaning: infinite and vicious regresses which they propose to cast out, substituting another regressive notion, such as mutual knowledge, instead; raising somewhat the question why their regresses are good regresses and mine are bad ones. (Grice, 1989, 299)

⁶According to Schiffer, once upon a time he was alone with Grice in an office in Berkeley and Grice agreed with him.

References

- Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.
- Clark, H. H. and Carlton, T. B. (1992). Context for comprehension. In Clark, H. H., editor, *Arenas of Language Use*, chapter 2, pages 60–77. University of Chicago Press and Center for the Study of Language and Information.
- Clark, H. H. and Marshall, C. R. (1992). Definite reference and mutual knowledge. In Clark, H. H., editor, *Arenas of Language Use*, chapter 1, pages 9–59. University of Chicago Press and Center for the Study of Language and Information.
- Elga, A. and Rayo, A. (2015). Fragmentation and information access. *Unpublished Manuscript*.
- Grice, P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass.
- Lederman, H. (2015). Uncommon knowledge. *Unpublished Manuscript*.
- Lewis, D. K. (1975). Languages and language. *Minnesota Studies in the Philosophy of Science*, VII:3–35.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15:509–559.
- Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.
- Schiffer, S. (1982). Intention-based semantics. *Notre Dame Journal of Formal Logic*, 23(2):119–156.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford, 2 edition.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics* 9, pages 315–332. Academic Press, New York.
- Stalnaker, R. (1984). *Inquiry*. MIT Press, Cambridge, Mass.
- Stalnaker, R. (1999). *Context and Content*. Oxford University Press.
- Stalnaker, R. (2014). *Context*. Context and Content. Oxford University Press, Oxford.
- Thomason, R. H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip. R. Cohen, J. M. and Pollack, M. E., editors, *Intentions in Communication*, chapter 16, pages 325–364. MIT Press, Cambridge, Mass.
- Yablo, S. (2014). *Aboutness*. Princeton University Press.