

Context as Knowledge

Torfinn Huvenes
Umeå University

Andreas Stokke
Swedish Collegium for Advanced Study
Uppsala University

--- forthcoming in *Mind and Language* ---

Abstract

According to an influential view, a context for a linguistic utterance is a body of shared information. The standard version of this view identifies this information with the information that is common ground in the conversation. Common ground information may be suited to the role that contexts play in the theory of communication and speech acts. Yet it has been argued that common ground information is less suited to the role that contexts play in the theory of indexical and demonstrative reference. This paper explores an alternative view that identifies shared information with what is common knowledge among the participants. We argue this view of shared information avoids the problems for the common ground approach concerning reference while preserving its advantages in accounting for communication.

1 Introduction

Utterances take place in contexts. The extent to which we rely on context in order to communicate makes it central to theories of languages and communication to ask what contexts are. According to a well-established view, a context for a linguistic utterance consists of information that is shared among the conversational participants. This generic view leaves it open what kind of information is relevant. The influential tradition associated with the work of Stalnaker (1974, 1975, 1978, 1998, 2002) identifies the information that makes up a context as, roughly, the information that is taken for

granted by the parties to the exchange.¹ This kind of shared information is known as the *common ground* of the conversation.

Contexts play a role in theorizing about communication and speech acts. But contexts also play a role in theorizing about indexical and demonstrative reference. While common ground information may be suited to the former role, it has been argued that common ground information is less suited to play the role that contexts have in determining reference. There is therefore a need for considering alternative approaches to contexts *qua* bodies of shared information.

In this paper we want to explore a different way of identifying shared contextual information. According to this alternative, shared contextual information is understood as what is *common knowledge* among the conversational participants. What makes up a context for a linguistic utterance is the information that is commonly known by the participants in the conversation.

We suggest that thinking about contexts in terms of common knowledge has some advantages over the orthodox view according to which contexts are bodies of information that are taken for granted. We will argue that there are reasons to think that taking contexts to be bodies of common knowledge avoids the problems concerning reference determination, while providing a satisfactory way of modeling communication.

At the same time, we will see that there are some challenges to this way of understanding contexts. While we think that the most critical challenges can be met, our aim here is to map out one potential alternative to the common ground picture of shared contextual information. We leave it open whether there are better alternatives, or indeed whether the general picture of contexts in terms of shared information should be abandoned.

Section 2 outlines the view that a context consists of the information that is common ground, along with some of its advantages and disadvantages. In Section 3 we present the alternative view that a context consists of information that is common knowledge. Section 4 addresses some objections.

2 Common Ground

For the purpose of our discussion, we rely on Stalnaker's (2002, 716) characterization of the common ground of a discourse. According to this picture, common ground information is defined in terms of two non-factive attitudes, acceptance and belief:

Common Ground

¹ Other theories of contexts in terms of information are found in, among others, Karttunen (1974), Gazdar (1979). Similarly, some versions of dynamic semantics assume that contexts consist of information, see e.g. Heim (1983), Veltman (1996), Groenendijk, Stokhof, and Veltman (1997). See e.g. Hunter (2013) for a discussion of indexicals in such a framework.

A proposition p is common ground in a group if and only if everyone in the group accepts that p , and everyone believes that everyone accepts that p , and everyone believes that everyone believes that everyone accepts that p , etc.

What is common ground in this sense corresponds to what is taken for granted, or presupposed. Accordingly, when we speak what is common ground in a conversation, we mean the presuppositions shared by the participants, unless it is made explicit that what is at stake is what is presupposed by a particular person.

The notion of common ground above represents contextual, shared information as a set of propositions. Yet it is sometimes convenient to represent shared information as a set of possible worlds, in particular, the set of possible worlds - called the *context set* - that are compatible with the information that is common ground. We follow Stalnaker in characterizing the context set of a discourse as follows.

Context Set

The context set of a conversation is the set of possible worlds w such that w is compatible with the information in the common ground.

Equivalently, the context set is the set of worlds w such that for all propositions p in the common ground, p is true in w . In what follows, we will sometimes speak of shared information as represented by the common ground and sometimes as represented by the context set.

2.1 Communication and Common Ground

The motivations for thinking about contexts in terms of common ground information are well-known. Stalnaker has emphasized repeatedly that contexts have two roles to play in theorizing about context-sensitive languages, and that the common ground is suited to play both of these roles:

[...] it is both the object on which speech acts act and the source of the information relative to which speech acts are interpreted. [...] So both of the roles that contexts play require that they include a body of information: context-dependence means dependence on certain facts, but the facts must be available, or presumed to be available, to the participants in the conversation. So I propose to identify a context (at a particular point in a discourse) with the body of information that is presumed, at that point, to be common to the participants in the conversation. (Stalnaker, 1998, 98)

So, on this picture communication proceeds against a background of shared information that serves to support interpretation of utterances. In turn, speech acts typically aim at updating this commonly shared information.

There are two features of the common ground, as characterized above, that are particularly important for our purposes. Both of these flow from the characterization

of common ground information in terms of acceptance. The attitude of acceptance is a non-factive attitude that is weaker than belief. Hence, the characterization above implies that, first, false propositions can be common ground. And second, a proposition can be common ground even if it is not believed. These features make it possible to account for the observation that an assertion does not have to be true, or even believed, to update the shared information in the conversation. Communication can still proceed against the backdrop of information that is false or not believed.²

Imagine that Harry asserts that Dale is in Washington. Yet Dale is not in Washington, and none of the conversational participants believe that he is. Even so, the conversationalists are willing to go along with Harry's assertion and accept that Dale is in Washington for the purpose of the conversation. (Perhaps they want to humor him, or perhaps it is a convenient short-cut to talking about what they are more interested in.) In this case it becomes common ground that Dale is in Washington. As such, the assertion has an effect on the interpretation of subsequent speech acts. For instance, following Harry's assertion, Lucy can felicitously ask why Dale is in Washington, thereby presupposing that he is. The presupposition is satisfied because it is common ground that Dale is in Washington. Correspondingly, every possible world in the context set is a possible world in which Dale is in Washington.

For the purpose of the following discussion, we are not going to challenge the claim that the common ground is suited to play a role in the theory of communication and speech acts, as illustrated above. However, the common ground approach to shared information raises questions about the role that contexts play in the theory of indexical and demonstrative reference.

2.2 Reference and Common Ground

We can distinguish two main ways of thinking about the common ground framework for theorizing about context-sensitive languages.³ On one approach, while common ground plays the central role in communication, indexical and demonstrative reference is determined by other aspects of contexts. In particular, reference will be seen as determined by facts about the utterance situation, such as who is speaking where and when. This approach has been advocated by Stalnaker (2014):

When I said to O'Leary, "You are a fool", the content of my speech act depended on the fact that O'Leary was my addressee, a fact that (like all facts) is determined by the K-context in which I am speaking. But if O'Leary does not realize that I am talking to him, then communication will not be successful. [...] So the account of context we need for our background story must distinguish a body of information that is available, or presumed to be available, as a resource for communication. (Stalnaker, 2014, 24)

² See Stalnaker (2002, 717–718).

³ Thanks to an anonymous reviewer for pressing this point.

By a “K-context,” Stalnaker has in mind the context *qua* collection of linguistically relevant facts about the utterance situation, such as who is speaking, where and when. This contrasts to the notion of context *qua* shared information. So, on this view, reference is determined by this kind of factual notion of context, while shared information is appealed to in order to understand various communicative phenomena.⁴

The other way of understanding the role of common ground in a theory of context-sensitivity is more unified. It sees shared information, such as common ground, as playing the central role both in communication and reference determination. As we noted earlier, in other work Stalnaker sometimes appears to endorse this more unified approach, and others have likewise understood the proposal this way.⁵ It is fair to say, however, that the most standard view embraces some version of the first of these.

For the purpose of the present discussion, it does not matter whether Stalnaker or anyone else has defended the view that a notion of context *qua* shared information is sufficient for theorizing about both communication and reference. Our interest here is in the question of whether such a unified approach can be made viable. What is important for our purposes, is that there are reasons to think that it cannot be made viable if the relevant shared information is the information that is common ground. In particular, it has been argued that common ground cannot play the role that context has in determining reference. Here is an example from Huvenes and Stokke (2016):⁶

Castor and Pollux

Castor and Pollux are identical twins, and are virtually indistinguishable to the naked eye and ear. In school Castor was a model student, while Pollux got into trouble and was ultimately expelled. Their parents managed to keep this event secret from people, and the twins changed schools so that after a few years their parents and themselves were the only ones to know about Pollux’s tainted academic history. Later in life, a series of physical accidents and psychological complications arose, and as a result people got the identities of the twins mixed up so that everyone, including themselves, came to believe that Castor is Pollux, and that Pollux is Castor. One day sitting in a café with some friends, Castor feels the urge to come clean, and after taking a deep breath, says:

(1) I was expelled from school.

Huvenes and Stokke make three observations about this example. First, the occurrence of *I* in (1) refers to Castor, the person who is in fact speaking. Second, (1) is false in this

⁴ See also Stalnaker (1999, 6-10).

⁵ See e.g. Stalnaker (1998, 101), Glanzberg (2002, 337), Szabó (2006, 384, fn. 59), and Gauker (2008, 194-196) also interpret Stalnaker as endorsing a unified view.

⁶ See e.g. Kaplan (1978) and Gauker (2008, 194-196) for similar examples and arguments involving other indexicals and demonstratives.

context and expresses the false proposition that Castor was expelled from school. These two observations are naturally related. If the occurrence of *I* in (1) had referred to Pollux, and not to Castor, (1) would have expressed the true proposition that Pollux was expelled from school, and not the false proposition that Castor was expelled from school. The third observation is that it is common ground in this conversation that Pollux is the speaker, or at least there is a natural way of fleshing out the example along these lines. After all, everyone believes that Pollux is the speaker and there is no reason to think that they would not be happy to take that for granted for the purpose of the conversation.

If the reference of *I* is seen as determined by the information that is common ground in the conversation, it is difficult to respect all of these observations. A natural suggestion given the identification of the context with common ground is as follows:

- (2) The referent of an occurrence *i* of *I* relative to a context set *S* is *x* if and only if
(i) *S* is non-empty and (ii) *x* is the speaker of *i* in every possible world in *S*.

Insofar as it is common ground that Pollux is the speaker in the example above, every world in the context is a possible world in which Pollux is the speaker. So, given (2), the relevant occurrence of *I* will refer to Pollux, and (1) will express the true proposition that Pollux was expelled from school. This conflicts with the observations about the example.

It might be pointed out, quite reasonably, that it may also be common ground that Castor is the speaker. After all, everyone believes of the man speaking right in front of them that he is the speaker. That man is Castor. Accordingly, this may equally be something that is accepted for the purpose of the conversation. But even if this is correct, it does not solve the problem as long as it is also common ground that Pollux is the speaker. In that case the common ground includes both the proposition that Castor is the speaker and the proposition that Pollux is the speaker. Since it is impossible that these two propositions are both true, this kind of common ground determines an empty context set. Consequently, given (2), this will be a case of reference failure.⁷ Yet this conflicts with the observation that *I* in (1) refers to Castor. As long as it is common ground that Pollux is the speaker, it is difficult to respect this observation.

There are several ways to respond to this problem. A lot more can be said on behalf of the common ground view. Some might be inclined to look for a way to solve the problem while still identifying the context with the common ground. Yet another reaction is to conclude that it is a mistake to characterize contexts wholly in terms of shared information. Reference is just not determined by shared information. Instead,

⁷ One could adopt an alternative way of determining reference, such that both Castor and Pollux are referred to when it is both common ground that Castor is the speaker and common ground that Pollux is the speaker. The fact that this also makes Pollux a referent need not be problematic in and of itself. However, Huvenes and Stokke (2016, p. 309) argue that this leads to incorrect predictions about the truth-value of (1).

one might argue that it is also necessary to include further facts about the utterance situation, such as who is speaking, where, and when.

While there might be merits to these ideas, we want to pursue a different approach. This involves holding on to the original idea of thinking of contexts in terms of shared information, while rejecting the specific suggestion to identify contexts with information that is common ground. Instead, we propose to characterize shared information in terms of knowledge.

3 Common Knowledge

We propose to identify contexts with information that is common knowledge in the conversation. As usual, we understand common knowledge as follows:

Common Knowledge

A proposition p is common knowledge in a group if and only if everyone in the group knows that p , everyone knows that everyone knows that p , everyone knows that everyone knows that everyone knows that p , etc.

As for common ground information, it will sometimes be convenient to represent this information as a set of possible worlds, in particular, the set of worlds that are compatible with the common knowledge of the participants. We think of this as the *knowledge set* of the conversation:

Knowledge Set

The knowledge set of a conversation is the set of possible worlds w such that w is compatible with the information that is common knowledge.

The knowledge set is more formally defined as the set of worlds w such that all the propositions that are common knowledge are true in w .

Like common ground, common knowledge is defined in terms of iterated attitudes. While this feature of common ground and common knowledge is controversial, it will not play an important part of the following discussion.⁸ For the most part, what matters is whether everyone in the conversation knows something, not whether they know that everyone knows it, and so forth. In particular, this is what will be important as far as the cases like that of Castor and Pollux is concerned.

The main difference between common ground and common knowledge *qua* bodies of information is that the latter is determined by a factive attitude. Only truths can be known, and hence only truths can be common knowledge. In turn, this means

⁸ Whether or not to characterize the shared information in terms of iterated attitudes, is not a trivial choice. For some discussion, see e.g. Hawthorne and Magidor (2009), Stalnaker (2009), and Almotahari and Glick (2011). A thorough treatment of these issues is beyond the scope of the present discussion. However, none of our arguments depend on this choice. For the sake of simplicity, we therefore focus on common knowledge throughout.

that the actual world is included in the knowledge set of any conversation. Hence, the knowledge set is always non-empty, since at least the actual world is always compatible with what is known. Correspondingly, if two propositions cannot both be true, they cannot both be common knowledge, and hence there are no inconsistent bodies of common knowledge.⁹

We will argue that the factivity of knowledge makes common knowledge better suited to theorizing about indexical and demonstrative reference. Furthermore, we will argue that this does not have to come at the expense of compromising the role of context as shared information in the theory of communication and speech acts.

3.1 Reference and Common Knowledge

As we have seen, cases like that of Castor and Pollux present a problem for the common ground view. In particular, the problem arises because there is nothing that prevents false propositions, such as the proposition that Pollux is the speaker, from being common ground. This makes it difficult to respect the observation that *I* in (1) refers to Castor. However, since knowledge is factive, the false proposition that Pollux is the speaker cannot be common knowledge. This is the first step towards ensuring that we make the right prediction about the case of Castor and Pollux.

Consider the following alternative rule for determining the reference of an occurrence of *I*:

- (3) The referent of an occurrence *i* of *I* relative to a knowledge set *K* is *x* if and only if *x* is the speaker of *i* in every possible world in *K*.

Given this suggestion, in order to respect the observation that *I* in (1) refers to Castor, we need to argue that, in this situation, it is common knowledge that Castor is the speaker.

This may seem like a tall order. After all, it is part of the description of the case that everyone believes that Pollux is the speaker. Furthermore, while the audience members would assent to (4), they would not assent to (5).

- (4) Pollux is the speaker.
(5) Castor is the speaker.

However, as we noted earlier, it is plausible to think that the conversational participants also have a belief about the individual they can see and hear speaking right in front of them. They believe that that individual is the speaker of the relevant occurrence of *I*. This belief is about Castor. So, the proposition the participants believe is true at a world *w* if and only if Castor is the speaker of *i* in *w*.

⁹ If nothing is common knowledge, the knowledge set is the set of all possible worlds. Similarly, if nothing is common ground, the context set is the set of all possible worlds.

To be sure, the participants would not describe themselves as having a belief about Castor. Yet this is compatible with the observation just made. One way to put this is to say that while the participants do not think about Castor under the guise of *Castor*, they think about him under something like the guise of *that person speaking*, or something similar.

But moreover, that Castor is speaking is arguably not just something the participants believe, it is something they know. Seeing and hearing someone speak is generally sufficient for knowing that that person is speaking. This makes it plausible that the proposition that Castor is speaking is common knowledge.

Since knowledge is factive, it cannot be common knowledge that anyone else is the speaker. Even if the participants believe that Pollux is the speaker, this belief cannot amount to knowledge. So, the situation corresponding to the one where both the proposition that Castor is the speaker and the proposition the Pollux is the speaker are common ground is ruled out. The knowledge set is always non-empty.

Consequently, in the case of Castor's utterance, every world in the knowledge set is a world in which Castor is speaking. Hence, given (3), the relevant occurrence of *I* refers to Castor, and (1) expresses the false proposition that Castor was expelled from school. This is the result we want.

The case of Castor and Pollux is just one example. There are more complicated cases involving indexicals.¹⁰ However, a discussion of such cases is beyond the scope of this paper, the main purpose of which is to put forward a view about contexts. By itself, the common knowledge view is unlikely to shed much light on more complicated cases involving indexicals. In order to do so, it would require, at the very least, a more thorough discussion of who counts as the speaker and whether it is possible to refer to use *I* to refer to someone other than the speaker. That is a discussion that must be left for another occasion.

The case of Castor and Pollux still illustrates the general advantages of common knowledge, compared to common ground, when it comes to determining reference. The case of Castor and Pollux highlights the problems with letting false propositions play a role in determining reference. By theorizing in terms of common knowledge instead of common ground, we avoid these problems.

3.2 Communication and Common Knowledge

At this point one might worry that we have simply traded one problem for another. Even if common knowledge is more suited to determine indexical and demonstrative reference, it may seem clearly unable to play the role that contexts are supposed to play with respect to communication.

The original conception of contexts in terms of common ground was designed to allow for shared information that is neither true nor believed. One main motivation

¹⁰ A widely discussed case is the so-called "answering machine paradox", noted by Kaplan (1989b, 491), and discussed by Sidelle (1991), Predelli (1998), and many others. For an overview and discussion, see Cohen and Michaelson (2013). See e.g. also Sherman (2015) for other complicated cases.

for this was to acknowledge that assertions that are false or not believed may have an effect on shared information. However, since knowledge entails true belief, a proposition is common knowledge only if it is both true and believed by every conversational participant. This may seem problematic.

Consider our previous example in which Harry asserts that Dale is in Washington despite the fact that Dale is not in Washington and none of the conversational participants believe that he is. The falsity of the asserted proposition and the fact that the conversational participants do not believe it, prevents it from becoming common knowledge. Yet we want to accommodate the observation that if the assertion is accepted by the conversational participants, it still manages to change the context. This is revealed by the difference it makes to subsequent speech acts. As we noted, if everyone accepts Harry's assertion, Lucy can felicitously ask why Dale is in Washington, thereby presupposing that he is in Washington.

The general worry is that common knowledge cannot do the kind of work that common ground is doing in the theory of communication and speech acts. One might think that since the asserted proposition cannot become common knowledge, we will be left without an explanation of the effect on the shared information of the conversational participants.

However, an assertion can still have an effect on the context even if the asserted proposition does not become common knowledge. It is a familiar observation that an utterance can affect the context in more than one way. Stalnaker (1998) observes that an assertion typically has the effect that the proposition that the utterance was made is added to the shared information:

A successful statement will thus change the context in two different ways that need to be distinguished. First, the fact that the statement was made is information that is added to the context simply as a result of the manifestly observable event that it was made. Second (assuming the statement is not rejected), the content of the assertion will be added to the context. (Stalnaker, 1998, 103)

As emphasized by von Stechow (2008), this idea is central to the way in which theorists in Stalnaker's tradition accounts for key phenomena such as presupposition accommodation:

[...] an utterance will affect the common ground in two steps: (i) first, the fact that the utterance was made becomes common ground (and the participants may immediately draw inferences based on that fact, and perhaps adjust the common ground accordingly), (ii) then, assuming that the proper (implicit) negotiation has occurred, the asserted proposition is added to the common ground. (von Stechow, 2008, 143)

Yet there is no reason to think that utterances are limited to having an effect on the context in just two ways. In the case of assertion, even if the asserted proposition cannot become common knowledge, a variety of other propositions may still become common knowledge.

It is plausible to think that this is what happens when Harry asserts that Dale is in Washington, even though no one believes it. While the asserted proposition cannot become common knowledge, another proposition becomes common knowledge, namely the proposition that everyone in the conversation accepts that Dale is in Washington.

While we are not in a position to spell out a detailed account along these lines, there is reason to believe that it may be able to handle the same phenomena as Stalnaker's two-step conception of the contextual influence of utterances described above. In particular, if the other conversational participants had not accepted Harry's assertion, the proposition that it is accepted that Dale is in Washington would not have become common knowledge. In that case, perhaps only the proposition that the assertion was made would have become common knowledge. So, the common knowledge view is able to distinguish between cases in which what is asserted becomes part of what is subsequently taken for granted and cases in which that does not happen.

One might wonder whether this is still a unified account of contexts.¹¹ After all, there are two relevant bodies of information. There is the information that is common knowledge, but there is also the information that is commonly known to be accepted by everyone in the conversation. While this is correct, the account remains unified.

What is crucial is that facts about what is commonly known to be accepted are entailed by facts about what is commonly known. Once all the facts about what is commonly known are settled, so are all the facts about what is commonly known to be accepted. In that sense, there is no need to appeal to anything beyond the information that is common knowledge, and so there is still only one notion of context in play.

For instance, whether it is common knowledge that everyone accepts that Dale is in Washington is entailed by the facts about what is common knowledge in the conversation. If it is common knowledge that everyone accepts that Dale is in Washington, this simply means that the proposition that everyone accepts that Dale is in Washington is common knowledge.

It is instructive to see why someone who wants to defend the common ground view cannot make a similar move in order to deal with the case of Castor and Pollux. For instance, one idea is to appeal to the information that is both common ground and true. Since it is false that Pollux is the speaker, this would not be a part of the relevant information. However, this would no longer be a unified account. Facts about what is common ground and true are not entailed by facts about what is common ground. Even if all the facts about what is common ground are settled, something more is

¹¹ Thanks to an anonymous referee for helpful discussion.

needed in order to settle the facts about what is common ground and true.¹² For instance, even if it is common ground that Dale is in Washington, that still leaves it open whether the proposition is true or not.

3.3 Presupposition, Assertion, and Content

The ideas we have been developing above have further implications for the treatment of presuppositions and assertions. For instance, these ideas can be extended to account for speaker presuppositions. When Lucy asks why Dale is in Washington, she presupposes that he is in Washington. Determining whether this presupposition is satisfied is not simply a matter of determining whether it is common knowledge that Dale is in Washington. It is instead a matter of determining whether it is common knowledge that it is accepted that Dale is in Washington. In general, what matters when we are dealing with this sort of speaker presupposition is not whether the presupposition is common knowledge, but whether it is common knowledge that it is accepted. As we have seen, the common knowledge view distinguishes correctly between such cases.

A useful way to frame this point is in terms of what counts as a successful assertion. An assertion may count as successful even if the asserted proposition does not become common knowledge. It is sufficient that it becomes common knowledge that the asserted proposition is accepted. This allows us to say that Harry's assertion that Dale is in Washington is successful despite the fact that the asserted proposition cannot become common knowledge. That is compatible with there being a higher level of success that requires that the asserted proposition become common knowledge. Harry's assertion does not attain this level of success, but it still manages to change the context in a way that makes a difference to subsequent speech acts.

As this suggests, the common knowledge view does not make the same predictions as the original common ground view of communication in all cases.¹³ There can be cases in which a proposition p is common ground, but it is not common knowledge that p is accepted – perhaps because the relevant beliefs do not amount to knowledge. We take it that such situations are, in some sense, defective. In particular, one can view such cases as examples of a purely epistemic defect. However, the common knowledge view goes further by taking such cases to involve a communicatively relevant defect.

¹² The point can also be made in terms of supervenience. Facts about what is commonly known to be accepted supervene on facts about what is commonly known. There can be no difference in facts about what is commonly known to be accepted without a difference in facts about what is commonly known. However, facts about what is common ground and true do not supervene on facts about common ground. There can be a difference in facts about what is common ground and true without a difference in facts about what is common ground.

¹³ It is worth noting that whenever it is common knowledge that a proposition p is common ground, then the factivity of knowledge ensures that p is common ground. That means that the common knowledge view and the common ground view will make similar predictions in a wide range of cases. It is only when a proposition is common ground, but it is not common knowledge that it is common ground, that the views make different predictions about communication.

Relatedly, it is important to emphasize that the common knowledge view allows for a distinction between, on the one hand, the proposition communicated by an utterance, and on the other hand, the proposition expressed or truth-conditional content. For example, in the Castor and Pollux case, it is plausible to think that, as a result of Castor's utterance of (1), it becomes part of what is accepted in the conversation that Pollux was expelled from school. Indeed, it becomes common knowledge that this is accepted. The point we have been emphasizing is that this is not the proposition that is expressed by (1) in this context. The proposition expressed by (1) is the proposition that Castor was expelled from school.

This kind of distinction extends to non-declarative utterances, as well. For example, suppose that one of the interlocutors in the Castor and Pollux example, asks Castor the question (6) after he has made his confession.

(6) Why were you expelled from school?

According to the current proposal, *you* in (6) refers to Castor. Briefly, since everyone knows that the person addressed is the person right in front of them, and since that person happens to be Castor, the utterance of (6) amounts to asking why Castor was expelled from school.

According to standard views, this question presupposes that Castor was expelled from school.¹⁴ In other words, (6) suffers from presupposition failure. This corresponds to the sense that (6), as addressed to Castor, does not have an answer, since Castor was not expelled from school. At the same time, as we said earlier, we take it to be plausible that it is common knowledge that it is accepted that Pollux is the speaker. Hence, we can account for the accompanying sense that the question will be interpreted as asking why Pollux was expelled from school.

While there might be other obstacles when attempting to further develop the view, we think that, given the sketch we have given above, it is plausible that common knowledge can play a suitable role in the theory of communication and speech acts.

4 Objections and Replies

We have seen that the orthodox way of thinking about shared information as common ground encounters some problems concerning reference determination. Instead of abandoning the conception of contexts as consisting of shared information, we have been exploring an alternative approach according to which a context is a body of common knowledge. In this section we develop this proposal by addressing some potential problems.

¹⁴ See e.g. Levinson (1983, 184), Roberts (2012, 23).

4.1 Defective Contexts

One worry concerns situations in which inconsistent presuppositions result in what is typically called *defective contexts*.¹⁵ For instance, let us suppose that, for some reason, the participants have become confused and are taking for granted that Bob and Norma have the same height, even though they also accept that Bob is taller than Mike, and Mike is taller than Norma. The common ground framework has a straightforward way of marking such a context as defective. In particular, since the common ground in this case contains inconsistent propositions, no possible worlds are compatible with the common ground, and the context set is reduced to the empty set.

On the common knowledge approach, the defectiveness of such contexts cannot be represented in the same way. As noted earlier, no body of common knowledge contains inconsistent propositions. If a set of propositions is inconsistent, its members cannot all be true. Therefore, they cannot all be common knowledge. The factivity of knowledge likewise makes it impossible to eliminate the actual world from the knowledge set. The knowledge set can therefore never be empty. In other words, we cannot represent defective contexts as situations in which no worlds are compatible with the shared information.

However, it is unclear how much of a problem this is. While a body of common knowledge cannot contain inconsistent propositions, there can be situations in which an inconsistent set of propositions such that for every proposition in the set, it is common knowledge that the proposition is accepted. For instance, it might be common knowledge that it is accepted that Bob and Norma have the same height, while at the same time it is common knowledge that it is accepted that Bob is taller than Mike, and that Mike is taller than Norma. Plausibly, this is sufficient to mark the context as defective. In particular, it licenses inconsistent presuppositions. In such a context, speakers can felicitously presuppose both a proposition and one that is inconsistent with it. In that sense, the context is seen as defective in a similar way to how it would be represented by the common ground framework.

It is true that this makes the consequences of defectiveness less dramatic. Insofar as the knowledge set is non-empty, the context can still function in many ways. However, it is not obvious whether this is a problem. After all, in practice, communication need not be affected if the inconsistency is irrelevant to the topic of the conversation. The inconsistency might only become a problem that needs to be addressed once it becomes relevant to the conversation.

Hence, while common ground and common knowledge represent defectiveness differently, it is unclear that this is a point in favor of the common ground view.

¹⁵ A defective context need not be the result of inconsistent presuppositions. See e.g. Stalnaker (1978, 2002) for relevant discussion. However, for present purposes, this serves as a clear and straightforward illustration of defectiveness.

4.2 Unsuccessful Communication

We have emphasized that contexts play at least two theoretical roles. Contexts play a role in the theory of communication and speech acts and in the theory of indexical and demonstrative reference. According to the common knowledge view, the information that is common knowledge plays both of these roles. While it may appear beneficial to have such a unified account, there is a worry that it comes at cost. In particular, there is a risk that the common knowledge view lacks the resources to correctly distinguish between cases of successful and unsuccessful communication.¹⁶

This can be illustrated by considering the case of Castor and Pollux. We have already seen that the common knowledge view is able to deliver the correct verdict that the occurrence of *I* in (1) refers to Castor. However, even though this is a case of successful reference, it looks less like a case of successful communication. In particular, the conversational participants end up with conflicting beliefs about the referent of *I*. On the one hand, they correctly believe that Castor is the referent. On the other hand, they also falsely believe that Pollux is the referent. As a result, there is arguably a sense in which communication has not been successful.

On views according to which the different roles of contexts in the theory of communication and reference are played by different theoretical entities, it might be less surprising that referential success and communicative success can come apart in this way. However, the issue is whether the common knowledge view, while providing a unified account of the two roles, can account for the sense in which communication is unsuccessful in these cases. Indeed, if the conversational participants believe the speaker, it is plausible that it becomes common knowledge that Castor was expelled from school. In that sense, it might seem that focusing on common knowledge commits us to regarding this as a case of communicative success.

Yet the fact that something was communicated by an utterance does not rule out that there are other ways in which communication was not successful. We want to highlight two ways in which the common knowledge view can be said to count the Castor and Pollux example as a case of failed communication. First, even if the utterance succeeds in making it common knowledge that Castor was expelled from school, the conversational participants nevertheless falsely believe that Pollux is the referent of the relevant occurrence of *I*.¹⁷ We think this false belief about who is referred to captures at least one sense in which communication is unsuccessful.

Second, the case plausibly involves a defective context in the sense outlined in the previous section. As we noted in 2.2 above, it is reasonable to think that it is not just common ground that Pollux is the speaker. It is also common ground that Castor

¹⁶ Thanks to an anonymous reviewer for pressing this point.

¹⁷ There are also cases in which only some of the conversational participants are mistaken about the identity of the speaker. For instance, it might be that only the hearers are mistaken. The hearers believe that Pollux is the speaker, but the speaker, Castor, has no such false belief. In that case, only the hearers would have a false belief about the referent of *I*. However, this would arguably still be sufficient to ensure that there is sense in which communication was not successful. Similar considerations apply if only some of the hearers are mistaken.

is the speaker. In that case, it is plausible that both the proposition that it is accepted that Castor is the speaker and the proposition that it is accepted that Pollux is the speaker are common knowledge. This means that the context licenses inconsistent presuppositions. If this is right, the failure of communication is accompanied by a defect in the context. This is a further indication that there is a communicative problem.

There is a general question of how one should negotiate the trade-offs between more or less unified accounts of the different roles of contexts. However, we see no other way to approach this issue other than on a case by case basis. We take it that the common knowledge view has sufficient resources to explain why there appears to be a failure of communication in cases like that of Castor and Pollux. It is possible to accept that there are cases of referential success without communicative success even on a unified account of the different roles of contexts.

4.3 Acquaintance

The previous worries concerned the role of common knowledge in the theory of communication and speech acts. But there are also worries about the role of common knowledge in the theory of indexical and demonstrative reference. Even if theorizing in terms of common knowledge makes it possible to account for cases like that of Castor and Pollux, there are other cases of indexical and demonstrative reference that potentially look more problematic.

In general, on the view we have been developing, in order for the right propositions to be common knowledge, the conversational participants must be in a position to think about the relevant referents. For instance, in the case of Castor and Pollux, the conversational participants must be in a position to think about Castor in order to know that he is the speaker. It is plausible that the participants have beliefs about Castor *qua* speaker because they are witnessing him speaking.

The worry is that something of this sort is not always going to be the case, and in particular, in some cases, the conversational participants are not acquainted with the referent and are therefore unable to possess the relevant knowledge. The upshot is supposed to be that common knowledge is also ill-suited for theorizing about indexical and demonstrative reference. If that is right, it undermines the main motivation for identifying the context with the information that is common knowledge, as opposed to the information that is common ground. In turn, this would be therefore a challenge to the proposal to retain the information-based approach to contexts by shifting the focus from common ground to common knowledge.

Consider an example. While Laura is traveling, she calls her cousin Maddy. Maddy believes that Laura is in Smalltown, but as a matter of fact, Laura is in Tinyville. During the call, Laura says:

(7) It's raining here.

In this context, we take it that *here* refers to Tinyville and that (7) expresses the proposition that it is raining in Tinyville.

Further, on the common knowledge view, the reference of *here* would be determined by a rule like (8).

- (8) The referent of an occurrence *i* of *here* relative to a knowledge set *K* is *x* if and only if *x* is the location of the speaker of *i* in every possible world in *K*.

In turn, therefore, we need to argue that, in the context for (7), it is common knowledge that the speaker is in Tinyville. Consequently, Maddy must know that the speaker is in Tinyville. But it might be that Maddy has never visited or even heard of Tinyville. In that case, one might worry that she is not even in a position to form a belief about Tinyville.

It is natural to spell out this objection in terms of acquaintance. Following Hawthorne and Manley (2012), we can distinguish between causal and epistemic acquaintance relations. In the case of Castor and Pollux, the conversational participants can see and hear Castor. That puts them in both a causal and an epistemic relation to Castor. But Maddy is not in perceptual contact with Tinyville. And so the worry is that Maddy does not stand in either a causal or an epistemic relation to Tinyville.

There is a lot to say about this objection. Before outlining a response, some preliminary remarks are in order.¹⁸ For one thing, one might object to the general idea of an acquaintance constraint.¹⁹ While we do not intend the response to depend on a general rejection of acquaintance constraints, we want to flag that the idea of acquaintance is controversial.

Furthermore, acquaintance constraints come in various degrees of strength. For instance, Russell (1910) took acquaintance to amount to direct awareness of the object in question. This is generally interpreted as a very strong acquaintance constraint. It arguably even rules out acquaintance with ordinary, external objects in our immediate environment. It should come as no surprise if the view under discussion is incompatible with such a strong acquaintance constraint. That is a fate it shares with many other views. However, it may be compatible with weaker acquaintance constraints.

A comprehensive discussion of various acquaintance constraints is beyond the scope of the present discussion. For present purposes, we will be content to outline a general response. What matters on the common knowledge view is what Maddy is in a position to think and know after the assertion has been made. Before the assertion,

¹⁸ Acquaintance is generally construed as a constraint on singular thought, but it is not always clear what counts as a singular thought. However, we do not want to get distracted by issues like that. For present purposes, what matters is whether the content of Maddy's thought is true at a possible world if and only if the speaker is in Tinyville in that world. If that does not count as a singular thought, that just makes life easier for those who want to identify contexts with common knowledge.

¹⁹ See e.g. Jeshion (2010); Hawthorne and Manley (2012).

it may well be true that Maddy lacks causal and epistemic contact with Tinyville. But that changes once the assertion is made. Laura is already in causal contact with Tinyville and through her assertion the causal connection is extended to Maddy. Furthermore, Maddy is now in a position to gain information about Tinyville through Laura. That ensures that there is also a sense in which she is in epistemic contact with Tinyville. In this sense, Maddy can be said to be in both causal and epistemic contact with Tinyville.

This response rests on the assumption that acquaintance can be passed along from the speaker to her audience.²⁰ This is similar to Kripke's (1980) notion of a chain of communication for proper names. Just like a name is passed on from link to link, so is acquaintance. This strikes us as a reasonable assumption. An acquaintance constraint that does not allow such chains of acquaintance makes it difficult to see how we can think about many of the objects and individuals that we ordinarily take ourselves to be able to think about. For instance, it is natural to think that we are sometimes in a position to think about people who died before we were born. In such cases, our causal and epistemic contact with the object of our thought is likely to depend on a chain of acquaintance.

The goal is not to show that this is sufficient to satisfy any acquaintance constraint. We have already seen that there are acquaintance constraints that cannot be satisfied in this way. But insofar as the worry is that Maddy is not in causal or epistemic contact with Tinyville, there are good reasons to think that the challenge can be met.

Still, one can question how far this line should be pushed.²¹ Imagine that Ed is playing a game with Annie and Nadine. In this game Ed closes his eyes and either Annie or Nadine writes a note. Ed is then supposed to open his eyes and guess who wrote the note. Upon opening his eyes, Ed reads,

(9) Dear Ed, I have a crush on you.

In fact, it was Annie who wrote the note. The challenge raised by this kind of example is that, intuitively, it seems clear that *I* refers to Annie and yet it is not obvious that a view like ours can account for this fact about the case.

We want to mention two broad reactions someone sympathetic to our account might have to this kind of situation. First, our account predicts reference failure in (9) only if it is not the case that it is common knowledge, in the context in question, that Annie wrote the note (or is the speaker, in the relevant sense). So one reaction is to try to motivate that this information is common knowledge among all three. Given the approach we have been sketching above, an argument to this effect might proceed by trying to argue that Ed can think about Annie in this case.

²⁰ See e.g. (Recanati, 2012, ch. 3) for relevant discussion. Recanati talks about relations of mediated acquaintance that are established by communicative chains.

²¹ We owe the following example to an anonymous reviewer.

One might for instance suggest that Ed can think of whoever the speaker is, that she is the speaker. For instance, imagine that Ed says, “Whoever wrote the note has terrible handwriting.” In that case Nadine can reasonably say to Annie, “Ed thinks that your handwriting is terrible.” It is unclear how much weight one should put on these kinds of attributions.²² In particular, one might be reluctant to think that Ed has a singular thought about Annie. However, the common knowledge view need not require the thought to be a singular thought in any strong sense. What matters for securing reference is whether the content of Ed’s thought is true at a possible world if and only if Annie is the author of the note in that world. There are obviously remaining questions, but we mention this as one possible reaction.

Second, it is not unnatural to think that there is more than one context in play in this case. In particular, one main thought behind the approach to contexts as information is to take the relevant information to be shared or available to all the participants.²³ In this sense, we can plausibly distinguish between the context between Annie and Nadine and the context between Annie, Nadine, and Ed. In the former the information that Annie is the speaker is available as common knowledge. So, one potential response is that the sense that *I* refers to Annie in (9) is to be explained by pointing to the fact that there is one way of understanding the context for the note in which it is common knowledge that Annie is the speaker. Indeed, there is clearly something unusual about the conversational setting in this example. After all, the point of the game is that the intended recipient of the note does not have access to the information that is required for interpreting the utterance. Hence, we might expect that it can be hard to identify the context for the utterance.

4.4 Defeaters and Gettier Cases

Even if the conversational participants are in a position to think about the referent, it is a further question whether the conversational participants are in a sufficiently strong epistemic position to possess the relevant knowledge. There are many ways to press this objection. Here we will focus on two versions of the worry that are meant to apply to a wide range of cases. The first concerns defeaters and the second concerns Gettier cases (Gettier 1963).

In the case of Castor and Pollux the conversational participants can see and hear Castor speaking. This would usually be sufficient for them to know that Castor is the speaker. However, in this case the conversational participants also believe that Pollux is the speaker. For the sake of the argument, we may suppose that while the belief that Pollux is the speaker is false, it is reasonable in light of the evidence available to the conversational participants. Furthermore, the belief that Pollux is the speaker is incompatible with the belief that Castor is the speaker in the sense that it is impossible for both of them to be true. The worry is that their belief that Pollux is the

²² For discussion of similar cases, see e.g. Recanati (2012, 152), Blumberg & Holguín (forthcoming).

²³ See in particular Stalnaker (2002, 715-716).

speaker constitutes a defeater for their belief that Castor is the speaker.²⁴ In turn, the objection goes, since this defeater is undefeated, the conversational participants cannot be said to know that Castor is the speaker.

It may seem plausible that someone has a defeater for her belief when she reasonably believes something that is incompatible with it.²⁵ However, cases involving identity confusion introduce further complications. It is natural to say that Lois Lane knows that Superman can fly, despite believing that Clark Kent cannot fly. If that is true, we cannot treat the belief that Clark Kent cannot fly as a defeater for the belief that Superman can fly despite the fact that the beliefs are incompatible.

There are different ways of accommodating scenarios like this. One might for instance say that people have beliefs about objects or individuals under different guises or modes of presentation and that defeaters are sensitive to these guises. Lois Lane's belief that Clark Kent cannot fly does not constitute a defeater for her belief that Superman can fly because the guises are different.

This should make one wary of treating the belief that Pollux is the speaker as a defeater for the belief that Castor is the speaker. After all, this is also a case of identity confusion even if it involves two individuals rather than just one. In fact, whatever one says about the case of Superman and Clark Kent, there is no obvious reason why one could not tell a similar story about the case of Castor and Pollux. It is no more obvious to the conversational participants in the case of Castor and Pollux that their beliefs are incompatible than it is to Lois in the case of Superman and Clark Kent.

Cases involving identity confusion introduce complications. But as long as identity confusion is compatible with having knowledge about the relevant individuals, we see no obvious reason that could not be true in the case of Castor and Pollux as well.

Let us turn to Gettier cases. Some Gettier cases might seem to present a problem for the common knowledge view.²⁶ Let us suppose that Gordon is about to tell Albert and Denise who is going to be promoted. Albert and Denise have already spoken to Margaret, who told them that Denise will be promoted. They both take Margaret to be reliable, but unbeknownst to them, she simply made it all up. However, it turns out to be a lucky guess and in fact it is Denise who will be promoted. In this context Gordon says,

(10) You are going to be promoted.

It seems that Gordon has successfully referred to Denise. Yet it is unclear whether the common knowledge view predicts that. In order for Denise to be the referent of *you*, it has to be common knowledge that Denise is the addressee. While Albert and Denise

²⁴ More precisely, this would be a case in which a belief or mental state acts as a rebutting defeater, in the terminology of Pollock (1986).

²⁵ See e.g. Pollock (1986, 37–58).

²⁶ Thanks to an anonymous reviewer for mentioning this problem.

plausibly believe that Denise is the addressee, it is also plausible that they do not know this. After all, their belief is based on Margaret's unreliable testimony.

There are two main points that we want to make in response to this case. First, the common knowledge view predicts that Gordon has communicated something about Denise. There are a number of propositions that plausibly become common knowledge as a result of his assertion. This includes the proposition that everyone accepts that Gordon was addressing Denise and the proposition that everyone accepts that Denise is going to be promoted. In that sense, we take it that Gordon communicated something about Denise. What the common knowledge view does not predict is that Denise is the referent of *you*. Still, the question is how problematic that is given that the view can acknowledge that Gordon has communicated something about Denise.

Second, there is arguably something defective about Gordon's assertion. The case only presents a problem for the common knowledge view if Albert and Denise do not know that Denise is the addressee. That means that Gordon cannot provide any further evidence that would be sufficient for Albert and Denise to come know that Denise is the addressee. For instance, Gordon cannot be seen looking at or gesturing towards Denise. But then it seems that he has not done enough to make it possible to work out who the referent is.

This makes the predictions of the common knowledge view seem less strange. While Denise is not the referent of *you*, Gordon nevertheless succeeds in communicating something about Denise. We regard this as a reasonable way of making sense of the case. In particular, it accounts for both the sense of success and the sense of failure.

4.5 Demonstratives

So far we have focused on cases involving indexicals like *I*, *here*, and *you*. Yet there is a question of how the account can deal with cases involving demonstratives like *that* or *this*. In particular, another potential worry about the common knowledge view is that it makes it too easy to refer and communicate using demonstratives.²⁷ Indeed, given what we argued in section 4.3 above, it might be thought that our view makes it impossible to fail to refer or misunderstand what is being referred to in cases involving demonstratives.

The important point to emphasize is that, even though, as we have argued, it is easy to think about referents, it does not follow that reference itself is easy. Consider the following rule for determining the reference of occurrences of *that* on the common knowledge view:

²⁷ Thanks to an anonymous reviewer for pressing this point.

- (11) The referent of an occurrence *i* of *that* relative to a knowledge set *K* is *x* if and only if *x* is the demonstrated object for *i* in every possible world in *K*.

Given (11), for an occurrence of *that* to refer to an individual it needs to be the case that all the participants know, and hence believe, that the individual is the demonstrated object for that occurrence. Consequently, on the common knowledge view, demonstrative reference requires that all the participants be able to think about the referent. In section 4.3 we argued that this condition is easily met.

At the same time, (11) does not specify what it takes for something to be the demonstrated object. What we have argued for here is not an account of how indexicals and demonstratives work, but an account of what a context is. The common knowledge view is neutral on specific views about what determines the reference of indexicals and demonstratives.

As such, (11) is compatible with different views about what secures demonstrative reference. On some views what a demonstrative refers to is determined by what the speaker demonstrates.²⁸ On other views what a demonstrative refers to is determined by what the speaker intends.²⁹ In terms of (11), these are views about what it takes to be the demonstrated object. For instance, if demonstrative reference is determined by what the speaker demonstrates, then (11) states that an occurrence of *that* refers to an individual if and only if all the participants know that the speaker demonstrated the individual. Among other things, this requires that the participants are able to think about the individual. If what we said in section 4.3 is correct, then there are no obvious reasons why this would not be relatively easy. But that is not all that is required. The conversational participants would have to know that the speaker demonstrated the individual. Since knowledge is factive, it would have to be true that the speaker demonstrated the individual. This is not a trivial requirement and opens up the possibility of referential failure and misunderstanding.³⁰ So far, this is similar to how indexicals work. The referent of *I* has to satisfy the relevant condition, presumably that of being the speaker. Similarly, the referent of *that* has to satisfy the condition of being the demonstrated object, however that is spelled out.

If one is still worried that the requirements are too easily met, there is nothing in the common knowledge view that precludes imposing further conditions on something being the demonstrated object. For instance, one might impose an additional uptake condition to the effect that any informed and reasonable hearer would recognize what is being referred to.³¹ Whether such a condition is motivated is beyond the scope of the present discussion. As we have already noted, the purpose of

²⁸ See e.g. Kaplan (1989b).

²⁹ See e.g. Kaplan (1989a), Stokke (2010), King (2014).

³⁰ Similar considerations apply to views on which demonstrative reference is determined by speaker intentions.

³¹ See e.g. Neale (2007), Stokke (2010), King (2014).

the present discussion is to advance a view about contexts. This is compatible with different views about demonstratives.

How easy or hard it is to refer and communicate using demonstratives depends on one's theory of demonstratives and not just one's theory of contexts. Hence, the common knowledge does not imply that it is impossible to fail to refer or misunderstand what is being referred to in cases involving demonstratives.

5 Conclusion

The approach to contexts as bodies of shared information is an attractive one. Yet the orthodox implementation of this view, which focuses on common ground, faces a challenge when it comes to playing the role that contexts need to play in the theory of demonstrative and indexical reference.

We have argued that the alternative view that contexts are bodies of common knowledge deserves to be taken seriously. By identifying context with information that is common knowledge we avoid the problems for the common ground view concerning reference determination. Furthermore, we have argued that it is possible to clear away the most obvious obstacles that would appear to prevent common knowledge from also playing a role in the theory of communication and speech acts.

References

- Almotahari, M., & Glick, E. (2011). Context, content, and epistemic transparency. *Mind*, 119, 1087-1105.
- Blumberg, K. & Holguín, B. (forthcoming). Ultra-Liberal Attitude Reports. *Philosophical Studies*.
- Cohen, J., & Michaelson, E. (2013). Indexicality and the Answering Machine Paradox. *Philosophy Compass*, 8, 580-592.
- von Fintel, K. (2008). What is presupposition accommodation, again? *Philosophical Perspectives*, 22, 137–170.
- Gauker, C. (2008). Against Accommodation. *Philosophical Perspectives*, 22, 171-205.
- Gazdar, G. (1979). *Pragmatics: Implicature, presupposition and logical form*. New York: Academic Press.
- Glanzberg, M. (2002). Context and Discourse. *Mind and Language*, 17(4), 333-375.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23, 121-123.
- Groenendijk, J., Stokhof, M., & Veltman, F. (1997). Coreference and modality. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 179–214). Oxford: Blackwell.
- Hawthorne, J., & Magidor, O. (2009). Assertion, context, and epistemic accessibility. *Mind*, 118, 377–397.
- Hawthorne, J., & Manley, D. (2012). *The reference book*. Oxford University Press.

- Heim, I. (1983). On the projection problem for presuppositions. In P. Portner & B. Partee (Eds.), *Formal semantics – the essential readings* (pp. 249–260). Oxford: Blackwell, 2002.
- Hunter, J. (2013). Presuppositional Indexicals. *Journal of Semantics*, 30, 381-421.
- Huvenes, T., & Stokke, A. (2016). Information centrism and the nature of contexts. *Australasian Journal of Philosophy*, 94, 301-314.
- Jeshion, R. (2010). Singular thought: Acquaintance, semantic instrumentalism, and cognitivism. In R. Jeshion (Ed.), *New essays on singular thought* (pp. 105–141). Oxford University Press.
- Kaplan, D. (1978). Dthat. *Syntax and Semantics*, 9, 221–243.
- Kaplan, D. (1989a). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565-614). Oxford and New York: Oxford University Press.
- Kaplan, D. (1989b). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford and New York: Oxford University Press.
- Karttunen, L. (1974). Presupposition and linguistic context. *Theoretical Linguistics*, 1, 181–194.
- King, J.C. (2014). The Metasemantics of Contextual Sensitivity. In A. Burgess and B. Sherman (Eds.), *Metasemantics: New Essays on the Foundations of Meaning* (pp. 97-118). Oxford and New York: Oxford University Press
- Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Neale, S. (2007). On location. In M. O'Rourke & C. Washington (Eds.), *Situating semantics: Essays on the philosophy of John Perry* (pp. 251–395). Cambridge, MA: MIT Press.
- Pollock, J. (1986). *Contemporary theories of knowledge*. Towota, NJ: Rowan and Littlefield.
- Predelli, S. (1998). I Am Not Here Now. *Analysis*, 58, 107-115.
- Recanati, F. (2012). *Mental files*. Oxford University Press.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5, 1-69.
- Russell, B. (1910). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society*, 11.
- Sherman, B. (2015). Constructing Contexts. *Ergo*, 2, 581-605.
- Sidelle, A. (1991). The Answering Machine Paradox. *Canadian Journal of Philosophy*, 21, 525–539.
- Stalnaker, R. (1974). Pragmatic presuppositions. In *Context and content* (pp. 47–62). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (1975). Indicative conditionals. In *Context and content* (pp. 63–77). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (1978). Assertion. In *Context and content* (pp. 78–95). Oxford: Oxford University Press, 1999.
- Stalnaker, R. (1998). On the representation of context. In *Context and content* (pp. 96–114). Oxford and New York: Oxford University Press, 1999.

- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Stalnaker, R. (2009). On Hawthorne and Magidor on assertion, context, and epistemic accessibility. *Mind*, 118, 399–409.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- Stokke, A. (2010). Intention-Sensitive Semantics. *Synthese*, 175, 383–404.
- Szabó, Z.G. (2006). The Distinction between Semantics and Pragmatics. *The Oxford Handbook of Philosophy of Language*, ed. E. Lepore and B. Smith. Oxford: Oxford University Press, 361–389.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25, 221–261.
- Stalnaker, R. (1978). Assertion. In *Context and content* (pp. 78–95). Oxford: Oxford University Press, 1999.
- Stalnaker, R. (1998). On the representation of context. In *Context and content* (pp. 96–114). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Stalnaker, R. (2009). On Hawthorne and Magidor on assertion, context, and epistemic accessibility. *Mind*, 118, 399–409.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- Stokke, A. (2010). Intention-Sensitive Semantics. *Synthese*, 175, 383–404.
- Szabó, Z.G. (2006). The Distinction between Semantics and Pragmatics. *The Oxford Handbook of Philosophy of Language*, ed. E. Lepore and B. Smith. Oxford: Oxford University Press, 361–389.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25, 221–261.