

A PUZZLE ABOUT CONTEXT AND COMMUNICATIVE ACTS

Daniel W. Harris

Abstract

A context-directed theory of communicative acts is one that thinks of a communicative act as a proposal to change the context in some way. I focus on three influential examples: Robert Stalnaker's theory of assertion, Craige Roberts' theory of questions, and Paul Portner's theory of directives. These theories distinguish different categories of communicative acts by distinguishing the components of context that they aim to change. I argue that the components of context they posit turn out not to be distinct after all, and that these theories therefore collapse the taxonomic distinctions that they set out to draw. Although it might be possible to avoid this problem by devising a more adequate theory of the nature of context, I argue that it should be taken as a reductio of context-directed theories.

I Communicative Acts

A communicative act is the speaker's contribution to a potential episode of communication. It is whatever it is that a speaker has to do, and that their addressee must correctly interpret, in order for communication to happen.¹ Suppose, for example, that Sam uses (1) to request that Ann buy him a drink.

(1) You should buy me a drink.

In order for Sam to thereby communicate with Ann, she must interpret him as performing a request rather than an observation or a prediction, as addressing her rather than someone else, as requesting that she buy him a drink now, rather than next month, and so on.

The central task of a theory of communicative acts is to tell us what it takes

¹ Communicative acts are most often referred to as 'speech acts' or, more specifically, 'illocutionary acts'. My terminology, which follows Bach and Harnish (1979), is meant to signal that I am not interested in *conventional* illocutionary acts, such as performing a ceremony or testifying in court, that are performable only against the background of social or institutional conventions.

to perform a communicative act of a given kind. A principled taxonomy of kinds of communicative act should follow from such a theory.

As (1) illustrates, the grammatical properties of a sentence normally underdetermine the communicative act one performs with it on an occasion of use. So, given that both options are grammatically available, what is it that makes Sam's act a request for Ann to buy him a drink today, rather than, for example, an observation that Bob is required by law to buy him a drink next week?

2 A Preview

My aim in this essay is to articulate a challenge to a genre of theories that purport to answer this question. I will call these theories *context-directed theories* of communicative acts. The central idea of these theories is that performing a communicative act is a matter of proposing a change to the conversation's context, which can be thought of as a body of information about the state of the conversation.² Communicative acts of different kinds, on this view, are proposals to update different components of the context: assertions add to the *common ground*, questions propose a new issue to be resolved (a *question under discussion*), and directives propose a task for the addressee to complete (by adding to their *to-do list*). I will spell out these ideas in detail in §§3–5. Along the way, I will contrast context-directed theories with what I call *addressee-directed theories*, such as the theory of Grice (1957; 1968; 1969), on which a communicative act is an attempt to change an addressee's private mental state, rather than the public state of the context.

One way for a theory of communicative acts to fail would be if it did not distinguish two categories of communicative act that ought to be distinguished. In the absence of some surprising argument for radically revising our pre-theoretic taxonomy, we should expect that an adequate theory will draw a principled distinction between assertions and both questions and directives, for example.

In what follows, I will argue that it is quite difficult for context-directed

2 Contexts of the kind at issue here sometimes go under other names, including 'conversational score', 'discourse context', 'information state', and so on. Early defenses of context-directed theories include Carlson (1982); Cohen and Perrault (1979); Gazdar (1981); Hamblin (1971); Heim (1982, 1983); Kamp (1981); Lewis (1979); Stalnaker (1978). The family has now grown too large to comprehensively cite, though some notable recent additions include Beaver (2001); Condoravdi and Lauer (2012); Farkas and Bruce (2010); Ginzburg (2012); Gunlogson (2001); Murray (2014); Murray and Starr (MS); Portner (2004, 2007, 2012); Roberts (2004, 2012); Starr (ms, 2010, 2014); Thomason (1990); Veltman (1996); Yalcin (2007, 2012).

theories to avoid collapsing these importantly distinct categories. My argument depends on the premise, which I will defend in §§6–7, that components of context must be *psychologically reducible*, *public*, and *conversation-relative*.

In §8, I argue that components of context with these properties are wholly determined by facts about common ground. If so, I argue in §9, then questions and directives turn out to be assertions in disguise. Although there may be some ways to avoid this conclusion – see §10 – I am tempted to think that my argument should be taken as a *reductio* of context-directed theories, and that we should instead prefer the less trendy view, due to Paul Grice, that the essential aim of communicative act is to change the private mental state of an addressee, rather than a public state of the context.

3 Assertion as a Proposal to Add to Common Ground

The best known and most influential context-directed theory is Robert Stalnaker's account of assertion, which revolves around a component of context that Stalnaker calls *common ground* (sometimes abbreviated as 'CG').³ Informally, Stalnaker describes common ground as the body of information that the participants in a conversation are presupposing, or treating as true, for the purposes of the conversation. Formally, CG is modeled as a set of propositions – a set of sets of possible worlds. The set of worlds in the intersection of these propositions – i.e., the set of worlds compatible with the common ground – is called the *context set*.

Common ground is widely thought to play a variety of roles in conversation. If it is common ground that Sam is addressing Ann when he utters 'you', then his interlocutors will understand him to have said something about Ann. If it is common ground that Bob is the tallest person in the room, Sam will be understood to be referring to Bob when uttering 'the tallest person in the room'. These examples exemplify two roles that common ground is thought to play. First, it is what hearers rely on as a supplement when the information that is linguistically encoded in a speaker's utterance does not fully specify what they are saying. Second, because they know that hearers will rely on common ground in this way, speakers also rely on common ground when they design their utterances: if their aim is to be interpreted in thus-and-such way, they should produce an utterance that, when supplemented by information from

3 Stalnaker (1978, 1998, 2002, 2014).

the common ground, will be interpreted in thus-and-such way. I will refer to these as the *pragmatic* roles of common ground.

I have doubts about whether we should posit common ground to play these pragmatic roles.⁴ However, I will grant that common ground plays these roles for the sake of argument. My reason for doing so is that I wish to focus on a further role that common ground is purported to play – its role as the target of assertions.

According to Stalnaker, “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” (2014, 39). Specifically, “assertion should be understood as a proposal to change the context by adding the content to the information presupposed” (1999, 10). “This is”, Stalnaker tells us, “an account of the force of an assertion” – i.e., a theory of the nature of this paradigmatic communicative act (1999, 10).

What does it mean to say that an assertion is a *proposal* to add a proposition to the common ground? In some of his work, Stalnaker takes this notion of proposal as a primitive. Elsewhere, he explicates it in terms of speakers’ intentions. “In my account of assertion,” he summarizes in a recent paper, “the idea was to explain speech act force in terms of the way that a speech act is intended by the speaker to change the context” (Stalnaker, 2018).

Thus explicated, Stalnaker’s theory of communicative acts belongs in an intentionalist tradition that begins in the work of Grice (1957; 1968; 1969). According to Grice, to perform a communicative act is to make an utterance with the intention of changing the mental state of one’s addressee, partly by way of revealing this intention to them.⁵ On a first pass, to assert that it is after midnight is to make an utterance with the intention of getting one’s addressee to believe that it is after midnight, in part by getting them to recognize this intention. Stalnaker’s view is quite similar, except that the role of the addressee’s beliefs is played by the conversation’s common ground.

Beliefs are *private* mental states, not in the sense that others can’t know about them, but merely in the sense that it is in the nature of such states to reside in the mind of a single individual. It is in the nature of common ground, by contrast, to be *public*, in the sense that the information in the common ground is pooled by and equally accessible to all of a conversation’s participants. (I’ll say more about just what this publicity amounts to in §§6–7.) The *essential aim* of

4 Some of my reasons are spelled out in ‘We Talk to People, Not Contexts’ (MS).

5 Grice uses different vocabulary: he talks about ‘utterer’s occasion meaning’ rather than communicative acts. My vocabulary here makes it easier to discuss non-assertoric acts, and to compare Grice’s theory with others’.

a communicative act is the aim that makes it the kind of communicative act that it is. For Grice, communicative acts essentially aim at changing addressees' private mental states. The essential aim of an assertion is to change the addressee's beliefs, for example. For Stalnaker, the essential aim of an assertion is to get one's message out into the open – to make it public – by putting it in the common ground. (Of course, as we will see, adding information to the common ground *entails* changing one's addressee's private mental states; but that is not *all* that it entails.)

What advantages are there to the idea that a communicative act is an attempt to make one's message public among the participants in a conversation, rather than merely an attempt to influence an addressee's private state of mind? Both theories enjoy some intuitive support: it is intuitive to think of communication as the transfer of information from one agent to another, but it is also intuitive to think of at least many conversations as processes of iterative information pooling. I am not optimistic that the matter can be settled by deciding which theory best matches our intuitive impression of conversation. The issue will have to be settled theoretically.

Stalnaker's context-directed theory has one apparent theoretical advantage, which is that it seems to explain the ways in which successful assertions influence what can happen later in the conversation.⁶ Recall that common ground is, according to Stalnaker and many others, the information that hearers rely on when they interpret speakers' utterances and that speakers rely on when designing utterances for their hearers. In general, we are entitled to take a proposition for granted in these ways after it has been successfully asserted in a conversation. Suppose that Sam asserts that Bob is taller than anyone else in the room and nobody else objects to this claim; henceforth in the same conversation, the description, 'the tallest person in the room', can be used to refer to Bob. Stalnaker has a simple explanation for this: interlocutors can felicitously presuppose whatever is in the common ground, and a successful assertion is one whose content has made it into the common ground. The aim of an assertion, it is tempting to conclude, is to add its content to the common ground.⁷

Stalnaker's account of assertion has served as an influential model for con-

6 Context-directed theories mostly haven't been directly compared to addressee-directed theories in the literature, and so the theoretical virtue of the former that I am about to present has not often been articulated as such. (One exception is Thomason 1990.) However, it is the best explanation I know of for the popularity of context-directed theories as opposed to addressee-directed theories.

7 Elsewhere (Harris, MS), I argue that this line of thought depends on running together several different ways in which a communicative act can be understood as succeeding. I will summarize some of these points in §10.

text-directed theories of other communicative acts. In general, the method suggested is to understand each kind of communicative act by positing a new component of contexts, part of whose theoretical role is to serve as the target of acts of the kind in question. I will sketch two of the most influential proposals of this kind in the next two sections: Craig Roberts' idea that asking a question is proposing a new *question under discussion*, and Paul Portner's idea that issuing a directive is proposing a new addition to the context's *to-do list*.

4 Questions as Proposals of new Questions Under Discussion

The intuitive idea behind Roberts' theory is that asking a question is proposing that it should become the new topic of conversation – the new issue that the interlocutors are working to resolve.⁸ The question that is operative in this way is what Roberts calls the *question under discussion*, or QUD. Like the common ground, the question under discussion is an element of the context around which a conversation unfolds: it changes as old questions are answered or new ones are proposed, and imposes constraints on how a well-ordered conversation can progress.

Again, this theory differs from Grice's view mainly in that Roberts takes the essential aim of asking a question to be a change to the public state of the context, whereas Grice thinks of questions as attempts to influence the addressee's plans about what to say. And again, each view matches our intuitions about some conversations: Grice's view makes sense of situations in which a speaker poses a question to only one of several interlocutors. Roberts' view naturally applies to conversations in which questions are raised for anyone to answer, as part of a process by which a group of agents are working together to pool their information.

Roberts' theory is theoretically motivated in a way that parallels what I have taken to be the main motivation for Stalnaker's view of assertion. She argues that a conversation's QUD plays several pragmatic roles, aside from being what one targets in asking a question. Successfully asking a question whose content is *q* normally has the effect that *q* will play these roles going forward. Roberts'

8 Although I will focus on a simplified version of the theory articulated by Roberts (2012; 2018), there are many precursors to, variations on, and further developments of her central ideas. See, e.g. Asher and Lascarides (2003); Carlson (1982); Ciardelli et al. (2013); Gazdar (1981); Ginzburg (1994, 1995a,b, 1996, 2012); Groenendijk and Roelofsen (2009); Hamblin (1971); Murray (2014); Staff (ms, 2010).

explanation of this fact is that the essential aim of asking a question is for its content to be the new QUD.

First, Roberts argues that the QUD plays a role in determining whether a subsequent communicative act will be judged *relevant*. This is why Sam's response to Ann would be felicitous in (2) but probably not in (3).

- (2) ANN: Who will drive Tony to the airport?
 SAM: I will drive him.
- (3) ANN: Who will drive Tony to the airport?
 SAM: # I would like a burrito.

Roberts predicts that an assertion is relevant just in case adding its content to the CG makes the CG newly incompatible with at least one answer to the QUD. A proposition that meets this criterion a *contextually relevant partial answer* to the QUD. (A contextually relevant *complete* answer to the QUD is a proposition that, together with the CG, rules out all but one element of the QUD.) The intuitive idea is that a relevant assertion is one that makes at least some progress toward the immediate conversational goal on which interlocutors have coordinated, which is to answer the QUD. The specific prediction that this account makes about (3) is that Sam's response will be irrelevant unless it is common ground, for some reason, that Sam's desire for a burrito, together with the common ground, rules out at least one previously live option about who will drive Tony to the airport. This could happen, for example, if it happens to be common ground that the best burritos are at the airport and Sam will want to go there to get what he wants. In the absence of such a scenario, Sam's response will seem infelicitous.

Roberts also argues that the QUD plays an essential role in determining when it is relevant to ask a question. Sam's response to Ann would be felicitous in (4), for example, but probably not in (5).

- (4) ANN: Who will drive Tony to the airport?
 SAM: Will you drive him?
- (5) ANN: Who will drive Tony to the airport?
 SAM: # Is it Tuesday already?

According to Roberts, asking a relevant question means asking a contextual *subquestion* of the QUD. In general, q_2 is a (contextual) subquestion of q_1 just in case any proposition that (contextually) entails a complete answer to q_1 also (contextually) entails a complete answer to q_2 . (If q_2 is a (contextual) subques-

tion of q_1 , then q_1 is a (contextual) *superquestion* of q_2 .) A corollary of these definitions is that, in completely answering a subquestion of q , one partially answers q . By answering a subquestion of the QUD, one can therefore make progress toward answering the QUD itself. This is what makes it relevant to ask a subquestion of the QUD. The intuitive idea here is that asking a question is often a way of pursuing a *strategy of inquiry*, whose goal is to break down a big question into more easily answerable chunks. In (4), Sam's response to Ann will seem relevant in any context in which it is not already common ground that Ann won't drive Tony. His response in (5) will seem relevant only in contexts in which it is common ground that the question of who will drive Tony turns on whether it is Tuesday.

Aside from this role in determining the relevance of communicative acts, it is widely accepted that QUDs also play a role in determining when it is felicitous to use prosodic focus. In response to (6), for example, (7) is a felicitous response, but (8) and (9) aren't.

- (6) Who will drive Tony to the airport?
- (7) [JIM]_f will drive Tony to the airport.
- (8) Jim will drive [TONY]_f to the airport.
- (9) Jim will drive Tony to the [AIRport]_f.

By contrast, (8) – but not (7) or (9) – is a felicitous response to ‘who will Jim drive to the airport?’, and (9) – but not (7) or (8) – is a felicitous response to ‘where will Jim drive Tony?’. On Roberts' view, these felicity conditions result because a declarative sentence with prosodic focus carries a presupposition that a congruent question is the QUD. In order for a declarative sentence φ to be congruent with a question q , the element in φ that is prosodically focused must correspond to the wh-phrase in an interrogative clause whose semantic value is q .

Do we need to posit contextual QUDs, as Roberts does, in order to explain which communicative acts count as relevant, to explain why certain uses of prosodic focus are felicitous? I have my doubts, but I will suppress those doubts here, and assume for the sake of argument that QUDs play these roles. Again, my focus will be on whether QUDs plays the conceptually distinct role of being what speakers must propose coordination on in order to ask a question.

5 Directives as Proposals to Add to the To-Do List

Directive acts, such as requests and commands, are attempts to prompt others to action. Grice captured this idea by arguing that the essential aim of directing someone to φ is to get them to form an intention to φ . Following the same general strategy set out by Stalnaker and Roberts, context-directed theorists have developed accounts of directives that parallel Grice's view, but with a newly posited component of the context playing the role that Grice assigns to the addressee's intentions.

The example I will discuss is due to Portner (2004; 2007), who posits a component of context that he calls the *to-do list* (TDL), which functions as a public record of the actions to which interlocutors are committed, for the purposes of the conversation.⁹ Portner captures this idea by modeling the to-do list as a function from interlocutors to sets of properties, each of which is something that an agent could instantiate by complying with a directive. Portner also argues that properties of this kind are the semantic values of imperative clauses (Portner, 2004; Zanuttini et al., 2012). A directive, Portner argues, is a proposal to add a new property to the addressee's section of the TDL. Whereas Grice would say that the essential aim of directing you to keep reading is to get you to form an intention to keep reading, Portner would say that its essential aim is to add the property of continuing to read to your section of our shared TDL.

Like Stalnaker and Roberts, Portner argues that his new theoretical posit plays a couple of conversational roles in addition to serving as the target of directives. And, again following the same pattern, the fact that issuing a directive will influence the conversation in these ways provides Portner with his strongest support for the idea that directives are proposals to update TDLs, rather than proposals to change addressees' private intentions.

First, Portner defines a precise way in which, he says, the state of the TDL, together with the state of the CG, determine which of interlocutors' actions will be deemed rational, for the purposes of the conversation. Normally, the following discourse would result in Bob's inaction being deemed irrational, for example:

- (10) ANN: Go outside immediately
BOB: Okay. [Bob doesn't move.]

9 Although I will focus on Portner's theory, roughly the same considerations I raise here will apply to a range of similar views proposed by others. See, e.g., Condoravdi and Lauer (2012); von Stechow and Iatridou (2017); Roberts (2004, 2015, 2018); Starr (ms, 2010).

Portner explains this by defending a formal implementation of the idea that, “to be judged rational an agent must attempt to make as many properties on her To-Do List true as possible” (2004, 243). Provided that it is common ground that going outside immediately requires moving, Bob’s behavior in (10) will be deemed irrational according to this standard.

Second, Portner argues that the TDL plays a role in determining how deontic modals are evaluated. It is often the case that a successful directive seems to “make true” the corresponding deontic modal. Imagine the following scene unfolding in a classroom, for example:

- (11) TEACHER, TO ANN: Go outside!
BOB, TO CAROL: Ann should go outside.

Portner argues that this phenomenon arises because the TDL influences the ordering source relative to which deontic modals are evaluated. According to a standard account of modals, Bob’s utterance in (11) is true just in case Ann goes outside in the worlds that are compatible with a contextually salient modal base and most highly ranked according to a contextually salient ordering source (Kratzer, 1981). According to Portner (2007), the contextual TDL influences the contextually available ordering source, so that worlds that realize more of the properties on the TDL will be more highly ranked. By issuing her directive, the teacher in (11) thus turns the contextually salient ordering source into one that assigns a high rank to worlds in which Ann goes outside, thereby giving rise to a context that makes Bob’s utterance true.

As in Stalnaker’s use of CG and Roberts’ use of QUDs, I think that there are reasons to be skeptical of this account. But once again, I will grant, for the sake of argument, that the TDL plays these pragmatic roles. My focus is on whether Portner is right that, given these assumptions, adding to the TDL should also be thought of as the essential aim of directive acts.

6 The Psychological Reducibility of Context

Let us assume that the context of a conversation consists of a common ground, a question under discussion, and a to do list, each of which plays the pragmatic roles set out in the last three sections. But now a question: what kinds of things are these posits? What are their natures? So far, I have described these posits in terms of their functional roles. But what are the realizers of these functional

roles? In virtue of what underlying facts about a conversation and its participants is its context in a given state at a given time?

The most worked out attempt to answer this question, and the one I will focus on here, is due to Stalnaker, who argues that facts about context reduce to facts about the propositional attitudes of the participants in a conversation. Call this approach *psychologism*.¹⁰ Stalnaker works out this idea in detail in his account of common ground. In his most recent work on the subject (2014), he argues that a proposition p is common ground just in case the interlocutors *commonly accept* p for the purposes of the conversation. A group of agents commonly accept p just in case each accepts p , each accepts that each of the others accepts p , and so on. To accept p , according to Stalnaker (1984, 79–84), is to treat p as true, either because one believes p , or because one bears some more provisional belief-like attitude to p . So, what's in the common ground is whatever the interlocutors treat as true, treat it as true that the others are treating as true, and so on, for the purposes of the conversation.

Thinking of context as psychologically reducible is a good idea for several reasons. It allows us to understand the posits of pragmatic theory in terms of propositional attitudes, which we have an independent grip on and need to posit for other purposes anyway. More importantly, psychologism allows us to explain why the state of the context influences interlocutors' actions. After all: we normally explain an agent's actions by appealing to their propositional attitudes. If facts about context reduce to facts about propositional attitudes, it is clear why changing the context will shape interlocutors' actions in predictable ways.

7 The Publicity and Conversation-Relativity of Context

Stalnaker's account also explains two other crucial features of common ground, which I will call its *publicity* and *conversation-relativity*.

Common ground is public in the sense that each of the interlocutors has

- 10 The main alternative to psychologism is conventionalism, according to which the state of the context is determined by the previous history of the conversation, together with the linguistic conventions that determine how conversations may unfold (DeVault and Stone, 2006; Lepore and Stone, 2015). The biggest disadvantage of conventionalism arises in situations when all participants in the conversation are simultaneously mistaken about the state of the context. In situations of this kind, we should expect their representations of the context (i.e., the context of psychologism), rather than the conventionally determined context, to play the relevant explanatory roles. For a more detailed version of this argument, see Harris (MS).

equal epistemic access to it. This is guaranteed by Stalnaker's definition, on which, if an interlocutor doesn't accept (or accept that the others accept, etc.) p , then p simply isn't common ground.

The publicity of CG is a property that it must have in order to play the pragmatic roles that Stalnaker takes it to play. If two people disagree about who the tallest person in the room is, or if one of them is wrong about who the other one thinks it is, then they will be unable to communicate when using the phrase, 'the tallest person in the room': one of them will design an utterance in order to refer to one person, and the other will interpret it as referring to someone else. The solution to this is for the CG to be public, for speakers to assume that hearers will interpret them by drawing on the common ground, and for hearers to assume that speakers have done so.

The same point applies to the components of context posited by Roberts and Portner. If the QUD were not public, then much infelicity and miscommunication would ensue, as there will be nothing to ensure that the interlocutors agree about which communicative acts are relevant, or about which prosodic focus is licensed. Similarly, if the TDL weren't public, then interlocutors could routinely be expected to miscommunicate with deontic modals and to disagree about what it would be rational for them to do, for the purposes of the conversation. Quite generally, facts about context must be public.

The conversation-relativity of common ground means that a proposition is common ground only relative to a particular conversation. In particular, a proposition may be common ground even though none of the interlocutors actually believes it, so long as they are treating it as true (and accepting that the others are doing the same, etc.) for the purposes of the conversation. This follows from Stalnaker's definition of acceptance: since provisionally treating p as true for some temporary purpose is a way of accepting p – for example, when one assumes p for the sake of argument – an agent may accept p while disbelieving p .

To see an illustration of this feature of common ground, imagine the following scenario. 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

believes that James is Eve's assistant, they do commonly accept this proposition, since they will continue to treat it as true, and to treat it as true that the others treat it as true (etc.), for the purposes of this conversation. Suppose that in this context, Felix utters the following sentence:

(12) Your assistant makes excellent cocktails.

In this context, we should expect that Felix has referred to James by uttering 'your assistant', and that the others will interpret him as doing so, despite the fact that none of them actually believes that James is Eve's assistant. It should therefore be clear that it is what the three of them commonly accept (for the purposes of the conversation) rather than what they believe (much less commonly believe) that plays the role of common ground.

The components of context posited by Roberts and Portner must be conversation relative in the same way. To see this, let's continue our spy story by supposing that the ostensible purpose of their meeting is to coordinate on a plan for Eve and James to deliver a bomb to Felix. And let us suppose that a detonation code is needed to activate bombs of this kind. Felix would need the code to go with his bomb. The conversation might proceed as follows.

(13) FELIX: What is the detonation code?

(14) EVE: The code is '42'.

Roberts' theory predicts that, in uttering (13), Felix's aim is to establish a new QUD, thus giving rise to the shared goal of adding an answer to the common ground. In uttering (24), Eve answers the question with a relevant assertion.

But keep in mind that neither Eve nor Felix has the goal of adding the *real* detonation code for bombs of this kind to the common ground. To do that would be to give away a state secret. Moreover, they both know this. And so in order for the QUD to be such that Eve's assertion is relevant, the QUD mustn't be constituted by the interlocutors' actual goals. It must depend on some conversation-relative substitute – the sort of thing that agents can adopt for some purpose, while privately aiming in the opposite direction. (I'll consider what state this might be in §8.)

The same point applies to the TDL. Suppose that their conversation proceeds as follows:

(15) FELIX: Deliver the bomb on Tuesday at 4:30.

(16) EVE: Very well. You must be there to pick it up.

- (17) JAMES, to eve: We should leave for to make the delivery by noon on Tuesday.

Portner's theory predicts that, with (15), Felix proposes to add the property of delivering the bomb on Tuesday at 4:30 to Eve's section of the to-do list. The first part of Eve's response ('very well') signals that Felix's proposal has been a success. According to Portner's theory, this new addition to the to-do list should have the effect of making true various deontic modals by raising to salience ordering sources that assign high rank to worlds in which Eve delivers the bomb on Tuesday at 4:30. This prediction is borne out by the felicity of (16)–(17), on the assumption that the highest-ranked, relevant worlds in which Eve delivers the bomb at the appointed time are also worlds in which Felix is there to pick it up at that time, and in which Eve and James leave to make the delivery by noon on Tuesday.

But remember: Eve and James don't *actually* desire (or plan) to go through with this delivery, and Felix knows this. So facts about the to-do list can't be constituted by interlocutors' *actual* plans or desires. They must instead depend on some conversation-relative substitute thereof. (I will consider what states could play this role in §8.)

Each of the components of context must therefore be both public and conversation relative in order to play the pragmatic roles that they are taken to play. Roberts and Portner seemingly agree. Roberts says that a conversation's QUD "commits [interlocutors] to a common goal, finding the answer" (Roberts, 2012, 5), but clarifies that this goal of finding the answer is to be understood in a conversation-relative way: "On the present view, it is the common ground, not the speaker, that's 'informed,' and it is mutual-belief behavior, and not knowledge, that's sought" (Roberts, 2012, 6, fn.7). Similarly, Portner says that "We may think of the Common Ground and To-Do List as being the public, or interactional, counterparts of the individual agents beliefs and desires" (Portner, 2004, 242).

8 The Psychological Reducibility of QUDs and TDLs?

Stalnaker has given us a detailed recipe for reducing common ground to agents' attitudes in a way that accounts for publicity and conversation-relativity. Although Roberts and Portner are committed to there being some such

account of the components of context that they posit, neither has attempted to articulate such an account. In this section I will explore the options for how such a reduction could be accomplished.

Roberts and Portner both describe the QUD and the TDL in terms of action-guiding mental states: goals, plans, or desires. So an initial idea might be to reduce the QUD and TDL to mental states of these kinds. I will focus on TDLs for now, though QUDs raise roughly analogous considerations that I will summarize below. According to Portner, the TDL is the “public and interactional” counterpart of desires, and so I will mainly consider that formulation, though analogous considerations apply to formulations in terms of plans or goals. Here’s a first shot at a reduction:

TDL1 For any agent A and a property φ , φ is on A ’s section of the TDL iff the interlocutors desire for A to instantiate φ .

This won’t do, since it captures neither the publicity nor the conversation-relativity of TDLs. So let’s try to capture publicity first. One way to copy Stalnaker’s strategy would be to build the TDL out of an intersubjectively iterated network of a single attitude – in this case, desires (or plans). This suggestion would work as follows.

TDL2 For any agent A and a property φ , φ is on A ’s section of the TDL iff:

- (i) Each interlocutor desires for A to instantiate φ
- (ii) Each of the interlocutors desires for each of the others to desire for A to instantiate φ
- ⋮

This account fails for several reasons. One is that it does not capture the conversation-relativity of the TDL – a point I’ll come back to below. But a more pressing problem is that this account doesn’t capture the publicity of the TDL either. The problem is that iterated desire doesn’t account for publicity in the same ways that iterated acceptance does. I can desire for you to desire φ , even if I accept that you don’t actually desire φ , and in this case I might not proceed as if you desire φ . (A related problem arises for plans: I can plan for you to plan to φ , even if I don’t accept that you yet plan to φ , in which case I won’t act as though you plan to φ .) What this shows, I think, is that the only attitudes that can be intersubjectively iterated in a way that results in publicity are acceptance states.

We might therefore try to define the to-do list as follows:

- TDL3 For any agent A and a property φ , φ is on A's section of the TDL iff:
- (i) A desires to instantiate φ
 - (ii) Each of the interlocutors accepts that A desires to instantiate φ
 - (iii) Each of the interlocutors accepts that each accepts that A desires to instantiate φ
 - ⋮

This account captures the publicity of the TDL: a property can't be on the TDL without playing the appropriate role in *all* of the interlocutors' psychological economies, by virtue of their intersubjectively iterated states of acceptance. The problem now is conversation-relativity. As I showed in the last section, none of the interlocutors need actually desire (or plan) to instantiate φ in order for φ to be on the to-do list. So, clause (i) makes TDL3 too strong.

I can think of two ways to relax this account in order to make the TDL appropriately conversation-relative. The first is to posit a new kind of conversation-relative attitude that stands to desire as acceptance stands to belief. This new desire-like attitude (or genre of attitudes) would play a functional role akin to desire, but it would be possible to bear this attitude provisionally and only for some purpose, even if one does not actually possess the corresponding desire, in the same way that one can accept something provisionally and only for some purpose, without actually believing it. To bear this attitude to some eventuality would be to treat it as desirable for some perhaps temporary purpose – e.g., for the purposes of the conversation. I think it is clear that, unlike acceptance, which is a natural enough idea, we lack anything like this provisional, desire-like concept in our folk psychology. I will call it *provisional desire*. (We can similarly discuss *provisional plans* and *provisional goals*.) With this concept in hand, we can make the following modification to our account.

- TDL4 For any agent A and a property φ , φ is on A's section of the TDL iff:
- (i) A provisionally desires to instantiate φ
 - (ii) Each of the interlocutors accepts that A (provisionally) desires to instantiate φ
 - (iii) Each of the interlocutors accepts that each accepts that A (provisionally) desires to instantiate φ
 - ⋮

The second way of weakening the desire requirement in TDL3 is to remove it altogether. Instead of A desiring to φ plus common acceptance that A desires to φ , we suffice with just the latter:

- TDL5 For any agent A and a property φ , φ is on A's section of the TDL iff:
- (i) Each of the interlocutors accepts that A desires to instantiate φ
 - (ii) Each of the interlocutors accepts that each accepts that A desires to instantiate φ
 - ⋮

How can we choose between TDL4 and TDL5? Each has some intuitive benefits. On one hand, it seems intuitive that the TDL wouldn't boil down to just a bunch of acceptance states, given that it is supposed to be action-motivating in a way that the common ground is not. On the other hand, the concept of provisional desire might be absent from folk psychology for a reason: can we really get a clear enough grip on it for our theoretical purposes? I think the answer to this question is unclear.

Ultimately, however, the way to answer this question is to work out which of these options best lines up with the functional role whose realizer we're searching for. And considerations of this kind give us a good reason to prefer TDL5 over TDL4. To see why, consider again the scenario from earlier involving the spies posing as arms dealers – Eve, James, and Felix – and the following variation on (I3)–(I4).

- (I8) FELIX, to james: Give me the detonation code.
- (I9) EVE: The code is in my briefcase. James, you should go and get it.
- (I20) JAMES: Very well. [He finds a slip of paper with '42' written on it and hands it to Felix.]

According to Portner's account, what's happened here is that Felix, in uttering (I8), has added the property of giving Felix the detonation code to James' part of the to-do list. This has a couple of effects. First, it raises to salience ordering sources that assign high rank to worlds in which James gives Felix the code. The first half of Eve's utterance makes it common ground that accomplishing this task requires retrieving the code from her briefcase. Together, these two communicative acts have the effect of making the second half of Eve's utterance, a deontic modal, seem true in the context. Second, Felix's addition to the to-do list has the effect of making James' actions appear rational if, given what is common ground, they bring it about that James will give Felix the code. This is borne out by James' action in (I20).

But remember: James does not actually desire (or plan) to give Felix the code. Does he provisionally desire (or provisionally plan) to do so? That is, does he treat it as desirable to give Felix the code, for the purposes of the conversation?

No, he does not. What he gives Felix is a fake code, and he, Eve, and Felix all privately realize this, even as they act as if James is giving them a real code. In fact, there may not even be a real detonation code. And even if there were such a code, James would not desire (or treat it as desirable, even temporarily) to give it to Felix. Provisionally desiring, understood in the way I've sketched, does not behave in the right way to ground facts about the to-do list, given the functional role that Portner ascribes to the latter, and so TDL4 won't work.

On the other hand, TDL5 does seem to do the job. Following Felix's utterance of (18), our spies do commonly accept, for the purposes of the conversation, that James desires (or plans) to give Felix the detonation code, and this state of common acceptance seems to play the right role in both making Eve's use of 'should' acceptable and in making James' act seem rational. Of course, action cannot be motivated by acceptance alone, any more than belief, and needs to be accompanied by some sort of desire or plan. But we can explain the preceding situation by saying that our spies all desire (or plan) to act as arms dealers in their position would act – i.e., to keep up their pretense. But this sort of desire is a private mental state, and won't show up on their TDLs.

I have been discussing the TDL, but an analogous line of reasoning applies to the QUD. I won't go through that reasoning in detail, but here it is in brief. Roberts (2012; 2018) argues that the QUD represents interlocutors' shared goal about the conversation's immediate future. However, as I showed in §7, in order to play the roles that Roberts attributes to it, the QUD must be public and conversation-relative. In order to account for publicity, the QUD must be grounded in what the interlocutors commonly accept about their goals. In order to be conversation-relative, the QUD must not depend on their actual goals, or even on what they are treating as their goals for the purposes of the conversation. After all, when Felix asks what the detonation code is by uttering (13), nobody actually adopts the goal of telling him the code, or even treats this as their goal for the purposes of the conversation. Rather, their goal is to go along with the pretense of arranging an arms deal which they all privately know to be fake. And so the following appears to be the best account of what grounds the QUD:

QUD q is the QUD of a conversation iff:

- (i) Each of the interlocutors accepts that it is the immediate goal of the conversation to resolve q
- (ii) Each of the interlocutors accepts that each accepts that it is the immediate goal of the conversation to resolve q

...

I can't claim that the line of reasoning I've pursued in this section is definitive. There may be some other way of reducing components of context to interlocutors' propositional attitudes. However, I find each of the steps I've gone through to be persuasive, and so I conclude that the QUD and TDL of a conversation reduce to interlocutors' propositional attitudes in much the same way as the CG. All three components of context are a matter of what is commonly accepted for the purposes of the conversation.

9 Collapse

If my reasoning up until now has been sound, then we've reached the puzzle of this paper's title, which goes as follows.

My conclusion in §8 was that the CG, the QUD, and the TDL all reduce to the same facts about interlocutors' propositional attitudes. For any proposition p , p is in CG just in case the interlocutors commonly accept p for the purposes of the conversation. For any question q , q is the QUD just in case the interlocutors commonly accept that it is their immediate conversational goal to resolve q . But let p be the proposition that it is the interlocutors' immediate conversational goal to resolve q . In that case, the fact that q is the QUD *just is* the fact that p is common ground.¹¹ Likewise, by parity of reasoning, for any agent A and property φ , the fact that φ is on A 's section of the TDL *just is* the fact that the interlocutors commonly accept that A desires (or plans) to instantiate φ . It follows that the state of a conversation's QUD and the TDL are wholly determined by the state of its CG (but not vice versa).

Context-directed theorists distinguish assertions from other communicative acts by the fact that other acts aim to change components of discourse context other than CG. But if all facts about these other components of context *just are* facts about CG, then changes to them are changes to CG as well. It follows that the other communicative acts in question are merely assertions in disguise. It follows from Portner's view that directing someone to φ has the same essential aim as asserting that they desire to φ , for example, and (21) has the same communicative-act potential as (22):

(21) Buy me a drink.

11 I am assuming what I take to be a relatively benign metaphysical principle: for any facts, α , β , and γ , if α reduces to γ and β reduces to γ , then α just is β . If the reader doesn't like my talk of reduction, they may feel free to substitute the relation of total grounding, for which the analogous principle seems equally benign.

- (22) You desire to buy me a drink.

And it follows from Roberts' view that asking someone a question is the same as informing them that it is the conversation's goal to resolve this question, so that (23) has the same communicative-act potential as (24):

- (23) Are you a doctor?
 (24) Our current conversational goal is to resolve the question of whether you are a doctor.

But this conclusion certainly isn't right. Although it may be possible to use declaratives of these kinds in order to perform the corresponding directive or act of questioning, the latter acts will be indirect. We can see this from the fact that they are cancellable. One can disambiguate (22) and (24) to make it clear that they are mere assertions as follows, for example.

- (25) You desire to buy me a drink. I'm not trying to get you to buy me a drink, I'm just stating an observation.
 (26) Our current conversational goal is to resolve the question of whether you are a doctor. I am not proposing this as a new goal, but merely reminding you of a fact.

In explaining differences between different kinds of communicative acts in terms of the different components of context they aim at changing, then, dynamic theorists unwittingly collapse the distinctions that they seek to explain.

10 Escape Routes?

How can this puzzle be resolved?

One option would be to bite the bullet, and to accept that questions and directives are really just assertions by another name. In particular, this strategy would involve admitting that the act of directing someone to finish their vegetables is the same act as asserting that someone desires to finish their vegetables – a conclusion that flies in the face of both common sense and the whole history of speech-act theory.

A second option would be to deny that contexts are psychologically reducible. This is certainly an option, and some theorists have explored it, albeit for

reasons different than the ones I am suggesting here (e.g. Devault and Stone 2006). The main challenge for such accounts is to say what context is, and why it plays the various roles that it is purported to play. As I argued in §§6–7, psychologism offers a neat and clear account of how and why context does what it does. As I hinted in footnote 10, and as I have argued in greater detail elsewhere, those who deny psychologism about context face serious challenges in explaining how it works (Harris, 2016, MS; Harris et al., 2018).

A third option would be to accept some version of psychologism about context, but to argue that there is some method of reducing the QUD and TDL of a conversation that is superior to anything I considered in §8. As I said at the end of that section, this may be possible but I know of no attempts. This essay can be read as a challenge to context-directed theorists to devise a workable alternative way of spelling out the nature of their posits.

A final option – and my preferred option – is to give up on context-directed theories altogether. My motivation for this conclusion depends on recognizing that context-directed theories ask the components of context that they posit to do two different kinds of things. First, context is made to play what I have called *pragmatic* roles: it is taken into account in the design and interpretation of utterances, it determines the relevance of communicative acts and the felicity of prosodic focus, and it helps determine the semantic values of deontic modals and the conditions under which interlocutors will deem an action rational. Second, context is the *target* of communicative acts – that which communicative acts essentially aim to change. If my argument in this essay has been sound, then contexts can't play both of these roles. I have used the pragmatic roles of context to pin down the functional properties of the various components of context, and then argued that if the context components have these functional properties, they collapse together – a conclusion that leads to the collapse of the different kinds of communicative acts, provided that communicative acts are individuated by the components of context that they target.

It might now be tempting to try to save context-directed theories by abandoning the idea that context plays the pragmatic roles that it is commonly taken to play, or by holding that we should posit contexts of different kinds to play the pragmatic roles and the role of communicative acts' target. But this response is undermined by the fact that the only real motivation for context-directed theories of communicative acts, as opposed to addressee-directed theories, depends on the idea that the context that serves as the target of communicative acts is the same context that plays the pragmatic roles. This is a point that I belabored in §§3–5. To sum up a point I made repeatedly there: the best argument for a context-directed theory of such-and-such kind of act is that

context plays thus-and-so pragmatic roles, and performing a such-and-such act reliably influences how context will do this job. If we give up the idea that contexts (and the same contexts, no less) play both roles, then this motivation for context-directed theories goes away.

I am therefore tempted to conclude that we *should* give up context-directed theories. We thus remain free to hold that context plays the various pragmatic roles assigned to it. And we can even go on believing that communicative acts typically have the kinds of effects on context that context-directed theorists take them to have. However, we should deny that these effects on context are the essential aims of communicative acts. Instead, we can go along with Grice in thinking that the essential aim of a communicative act is the effect that it is intended to have on the private mental state of its addressee; its effect on context is a further, extra-communicative aim that we often accomplish indirectly, by means of achieving our essential communicative aims.¹² Since there is presumably no concern about the concepts of belief, desire, and intention collapsing into one another, addressee-directed theories are not susceptible to the kind of collapse argument that I directed at context-directed theories in §9. Although there is much to be said about how addressee-directed theories should be understood, I won't attempt to articulate any of that, except to point out that such theories avoid the puzzle that has been my concern in this essay.

II Conclusion

Context-directed theories work by individuating categories of speech acts in terms of the components of context that they aim to change. I have argued that the best motivations for these theories depend on taking the components of context that they posit to play certain functional roles, and that if the components of context really do play these roles, they aren't really distinct components of context at all. But by collapsing the components of context that they posit, context-directed theories collapse the very distinctions between different categories of communicative acts that they have been formulated to explain in the first place.

I find it tempting, for this and other reasons (Harris, MS), to conclude that we should abandon context-directed theories. But even if context-directed theorists are not tempted by my argument to give up hope in their views, they

12 I defend this view in greater detail elsewhere (Harris, 2014, MS).

should at least be motivated to properly theorize the components of context that they posit in a way that gets around the objections that I have raised here. To do so would be to fill a gaping foundational lacuna in their theories – one that I suspect cannot be filled.

Acknowledgement

This essay has been improved by feedback from Kent Bach, Justin Bledin, Eliot Michaelson, Paul Portner, Craige Roberts, Elmar Unnsteinsson, and Seth Yalcin.

References

- Asher, N. and Lascardes, A. (2003). *Logics of Conversation*. Cambridge University Press, Cambridge, UK.
- Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.
- Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI, Stanford, CA.
- Carlson, L. (1982). *Dialogue Games: An Approach to Discourse Analysis*. Number 17 in Synthese Language Library. D. Reidel, Dordrecht, The Netherlands.
- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2013). Inquisitive semantics: a new notion of meaning. *Language and Linguistics Compass*, 7(9):459–476.
- Cohen, P. R. and Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212.
- Condoravdi, C. and Lauer, S. (2012). Imperatives: Meaning and illocutionary force. In Pinoñ, C., editor, *Empirical Issues in Syntax and Semantics 9: Papers From the Colloque de Syntaxe et S'emantique a` Paris 2011*, pages 37–58.
- DeVault, D. and Stone, M. (2006). Scorekeeping in an uncertain language game. In Schlangen, D. and Fernáandez, R., editors, *Brandial 2006: Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, volume 10, pages 139–146. SemDial, University of Potsdam.
- Farkas, D. F. and Bruce, K. B. (2010). On reacting to assertions and polar questions. *Journal of Semantics*, 27:81–118.
- von Fintel, K. and Iatridou, S. (2017). A modest proposal for the meaning of imperatives. In Arregui, A., Rivero, M., and Salanova, A. P., editors, *Modality Across Semantic Categories*, chapter 13, pages 288–319. Oxford University Press, Oxford, UK.
- Gazdar, G. (1981). Speech act assignment. In Joshi, A. K., Weber, B. L., and Sag, I. A., editors, *Elements of Discourse Understanding*, pages 64–83. Cambridge University

- Press, Cambridge, UK.
- Ginzburg, J. (1994). An update semantics for dialogue. In *First International Workshop on Computational Semantics (IWCS 1)*, pages 111–120.
- Ginzburg, J. (1995a). Resolving questions, i. *Linguistics and Philosophy*, 18(5):459–527.
- Ginzburg, J. (1995b). Resolving questions, ii. *Linguistics and Philosophy*, 18(6):567–609.
- Ginzburg, J. (1996). Dynamics and semantics of dialogue. In Seligman, J. and Westerstahl, D., editors, *Language, Logic, and Computation*. CSLI, Stanford, CA.
- Ginzburg, J. (2012). *The Interactive Stance: Meaning for Conversation*. Oxford University Press, Oxford, UK.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.
- Grice, H. P. (1968). Utterer's meaning, sentence-meaning, and word-meaning. *Foundations of Language*, 4(3):225–242.
- Grice, H. P. (1969). Utterer's meaning and intention. *The Philosophical Review*, 78(2):147–177.
- Groenendijk, J. and Roelofsen, F. (2009). Inquisitive semantics and pragmatics. In Larrazabal, J. and Zubeldia, L., editors, *Meaning, Content, and Argument: Proceedings of the ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric*, pages 41–72. University of the Basque Country Publication Service.
- Gunlogson, C. (2001). *True to Form: Rising and Falling Declaratives as Questions in English*. Routledge.
- Hamblin, C. L. (1971). Mathematical models of dialogue. *Theoria*, 37:130–155.
- Harris, D. W. (2014). *Speech Act Theoretic Semantics*. PhD Dissertation, City University of New York Graduate Center.
- Harris, D. W. (2016). Intentionalism versus the new conventionalism. *Croatian Journal of Philosophy*, 16(47):173–201.
- Harris, D. W. (MS). We talk to people, not contexts. *Unpublished Manuscript*.
- Harris, D. W., Fogal, D., and Moss, M. (2018). Speech acts: The contemporary theoretical landscape. In Fogal, D., Harris, D. W., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD Dissertation, University of Massachusetts at Amherst.
- Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In Bäuerle, R., Schwarze, C., and von Stechow, A., editors, *Meaning, Use and Interpretation of Language*, pages 164–189. de Gruyter, Berlin.
- Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language, Part 1*, volume 135, pages 277–322. Mathematical Centre Tracts, Amsterdam.
- Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. J. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter, Berlin.
- Lepore, E. and Stone, M. (2015). *Imagination and Convention*. Oxford University Press, Oxford.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Murray, S. E. (2014). Varieties of update. *Semantics and Pragmatics*, 7(2):1–53.
- Murray, S. E. and Starr, W. B. (MS). The structure of communicative acts. *Unpublished*

- Manuscript, Online at <https://www.dropbox.com/s/ctf8stqvlplrdqs/The-Structure-of-Communicative-Acts.pdf?dl=1>.*
- Portner, P. (2004). The semantics of imperatives within a theory of clause-types. In Watanabe, K. and Young, R., editors, *Proceedings of SALT 14*. CLC Publications.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15:351–383.
- Portner, P. (2012). Permission and choice. In Grewendorf, G. and Zimmermann, T., editors, *Discourse and Grammar: From Sentence Types to Lexical Categories*. Mouton de Gruyter.
- Roberts, C. (2004). Context in dynamic interpretation. In Horn, L. and Ward, G., editors, *The Handbook of Pragmatics*, pages 197–220. Blackwell.
- Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.
- Roberts, C. (2015). Conditional plans and imperatives: A semantics and pragmatics for imperative mood. In Brochhagen, T., Roelofsen, F., and Theiler, N., editors, *Proceedings of the 20th Amsterdam Colloquium*, pages 353–362.
- Roberts, C. (2018). Speech acts in discourse context. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*, Oxford, UK. Oxford University Press.
- 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. (1998). On the representation of context. *Journal of Logic, Language, and Information*, 7(1):3–19.
- Stalnaker, R. (1999). *Context and Content*. Oxford University Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, e25:701–721.
- Stalnaker, R. (2014). *Context*. Context and Content. Oxford University Press, Oxford.
- Stalnaker, R. (2018). Dynamic pragmatics. In Fogal, D., Harris, D. W., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press.
- Starr, W. (ms). A preference semantics for imperatives. *Unpublished Manuscript*.
- Starr, W. B. (2010). *Conditionals, Meaning and Mood*. PhD thesis, Rutgers University.
- Starr, W. B. (2014). Mood, force, and truth. *Protosociology*, 31:160–180.
- 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.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.
- Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160.
- Zanuttini, R., Pak, M., and Portner, P. (2012). A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language & Linguistic Theory*, 30(4):1231–1274.