

We Talk to People, not Contexts

Daniel W. Harris

Draft Completed June 16, 2017^{*}

A communicative act is whatever it is that a speaker has to do, and that an addressee has to correctly interpret, in order for successful communication to happen. Paradigm cases are assertions, questions, requests, and commands—the sorts of acts that one performs by using unembedded declarative, interrogative, and imperative sentences in a literal and direct way.¹ My aim here is to compare two families of theory of the nature of communicative acts, one trendy and the other a bit old fashioned. I will then defend the old-fashioned option.

The trendy theories that I will object to are what I'll call *context-directed* theories of communicative acts. According to theories of this kind, to perform a communicative act is to do something whose aim is to change the conversation's context in some way. For the purposes of these theories, which are usually elements in dynamic-semantic or dynamic-pragmatic models of conversation, the context of a conversation is specific kind of theoretical posit—a body of representational contents that is intersubjectively constructed out of the private mental states of the conversation's participants (who are also known as 'the interlocutors'). Stalnaker's (1978; 2014) influential theory of assertion has it, for example, that to assert p is to

^{*}I have received helpful feedback on earlier versions of some of the ideas presented here from Kent Bach, Anne Bezuidenhout, Justin Bledin, Liz Camp (and the students in her Spring 2017 graduate seminar at Rutgers), Kai von Fintel, Daniel Fogal, Ernie Lepore, Matt Mandelkern, Eliot Michaelson, Sarah Murray, Stephen Neale, Craige Roberts, Will Starr, Elmar Unnsteinsson, and Seth Yalcin. This draft is not final, and I would appreciate feedback, which can be emailed to me at danielwharris@gmail.com.

¹I have chosen to talk about 'communicative acts' rather than 'speech acts' both because I think that there may be non-linguistic communicative acts, and to focus attention away from conventional speech acts, which can be performed only against the background of social conventions or institutions. On the distinction between communicative and conventional acts, see Bach and Har-nish (1979).

act with the aim of adding p to the common ground, which is the set of propositions that the interlocutors commonly accept for the purposes of the conversation. Other kinds of communicative acts work the same way, except that they are essentially aimed at changing different components of the context.

The old-fashioned theories that I wish to defend are what I'll call *addressee-directed* theories of communicative acts. According to these theories, to perform a communicative act is to act with the aim of changing the addressee's mental state in some way. For example, according to Grice (1957; 1968; 1969), to assert p is to do something with the aim of getting one's addressee to form a belief in p , in part by getting them to recognize that this is one's aim. Communicative acts of other kinds work the same way, but are aimed at changing mental states other than beliefs.²

It should be stressed that addressee-directed theories are compatible with the existence of contexts, and with the idea that we sometimes aim to influence contexts with out communicative acts. However, if an addressee-directed theory is correct, then changing the context is, at best, an ancillary aim communicative acts, and one that needn't always be present. To change the context is something over and above the kind of success at which one must aim in order to communicate at all.

The debate between context-directed theories and addressee-directed theories can seem like an internecine one, since the two views have a lot in common. Both take the nature of a communicative act to boil down to its aim, which is to influence a body of representations. The two views differ only on the question of whether this body is a private state of an individual's mind, or a public and intersubjective state of the context. Other theories of communicative acts eschew aims altogether, instead understanding communicative acts in terms of the conventions or norms that govern them, or in terms of the mental states they express.³ For the purposes of this essay, I will ignore these other views, and pretend that context-directed theories and addressee-directed theories are the only games in town.

My plan is to defend addressee-directed views by giving examples of communication that they make better sense of in §2, and by carefully distinguishing a

²Grice calls assertions 'indicative-type utterances' and directives 'imperative-type utterances'. He initially (1957) takes directives to be constitutive aim to be producing action rather than an intention to act (see also (Schiffer, 1972, ch.4)), but changes his mind in later work (1968; 1969).

³For convention-based views, see Alston (2000); Austin (1962); Lepore and Stone (2015); Searle (1965, 1969); Searle and Vanderveken (1985). For norm-based views, see Brandom (1983, 1994, 2000); Kukla (2014); Kukla and Lance (2009); MacFarlane (2011, 2014); Williamson (1996, 2000). For the view that speech acts should be understood in terms of the psychological states they express, rather than the ones they aim to produce, see Bar-On (2013); Green (2007); Pagin (2011); Rosenthal (1986). For an overview of how the different kinds of theories of communicative acts relate to each other, see Harris et al. (2017).

variety of ways in which a communicative act can succeed in §3. In §4, I will try to pin down the sense in which context-directed theories could be thought of as the product of idealized models of conversation. First, in §1, I will make it as clear as possible what is at issue in the debate.

1 Context-Directed Theories of Communicative Acts

The central thesis of context-directed theories of communicative acts can be stated as follows.⁴

(CD) CONTEXT DIRECTEDNESS

The essential aim of an communicative act is to change the context in some way. Communicative acts of different kinds aim to change the context in different ways.

The best known theory of this kind is Robert Stalnaker's theory of assertion, which I sketched in the last section. However, others have extended Stalnaker's theory in a variety of ways by positing further components of context for other kinds of communicative acts to manipulate. For example, Roberts (2004; 2012; 2017) argues that the context includes a *question-under-discussion stack*, that the question at the top of the stack at a given moment is the immediate question under discussion (QUD), and that the aim of asking a question is to install a new QUD at the top of the stack.⁵ In a similar vein, Portner (2004; 2007) argues that the context includes a *to-do list* (TDL)—a public record of the practical commitments on which interlocutors have coordinated—and that the essential aim of performing a

⁴Influential early statements of this idea include Carlson (1982); Cohen and Perrault (1979); Gazdar (1981); Hamblin (1971); Heim (1982, 1983a); Kamp (1981); Lewis (1979); Stalnaker (1978). The idea is now too widespread to comprehensively cite, though some influential recent contributions that focus on the nature of communicative acts include Beaver (2001); Clark (1996); Condoravdi and Lauer (2012); Farkas and Bruce (2010); Ginzburg (2012); Gunlogson (2001); Murray (2014); Murray and Starr (MS); Portner (2004, 2007, 2012, 2017); Roberts (2004, 2012, 2017); Starr (ms, 2010, 2014); Thomason (1990); Veltman (1996); Yalcin (2007, 2012, 2017).

⁵I will focus on Roberts' influential version of these ideas because Roberts clearly signals her allegiance to a context-directed theory of speech acts as I am defining it here. However, several of the crucial aspects of Roberts' theory of questions can be found elsewhere as well: Carlson (1982); Ginzburg (1994, 1995a,b, 1996, 2012); Hamblin (1971); Murray (2014); Murray and Starr (2017); Portner (2004); Starr (ms, 2010).

directive speech act is to add a new commitment to the addressee's section of the TDL.⁶

There are several senses in which an act may have an aim. For example, some have held that the aim of a communicative act is a matter of the intention with which it is performed.⁷ Others have argued that the aims of at least some communicative acts should be understood as the effects that it is their function to produce, and that an act's function arises not from the speaker's intentions but processes akin to natural selection.⁸ Either of these conceptions of the aim of a communicative act may be put to work, or combined, in either a context-directed or addressee-directed theory. Although some context-directed theorists have openly embraced versions of intentionalism (e.g. Roberts 2012; 2017; Thomason 1990), others have mostly dodged the issue of where speech acts' aims come from. Stalnaker's way of doing this, which has caught on among some others, is to say that assertions are "proposals" to change the context, without saying anything further about what a "proposal" is.⁹ I will mostly focus on intentionalist views, mainly because they're more nuanced and better worked out at present. However, I think that the central points in this essay would translate to functionalist theories.

The word 'essential' in (CD) is crucial. Most communicative acts are aimed at many different outcomes, since they fit into the speaker's plans in a variety of ways. In telling you that I am tired, I may aim to share the information that I am tired, to motivate you to leave my apartment, to cause you to realize that you have overstayed your welcome, to deter you from doing so in the future, and to get to bed soon. Plausibly, only the first of these is the essential aim my assertion. (The second, but none of the rest, may be the essential aim of a further act of implicating that you should leave.) So, a theory that grounds the nature of communicative acts in their aims needs a way of distinguishing between an act's essential and an-

⁶For views that are similar enough to Portner's as to be interchangeable for present purposes, see Condoravdi and Lauer (2012); von Stechow and Iatridou (ms); Roberts (2004, 2015b, 2017); Starr (ms, 2010).

⁷Bach and Harnish (1979); Cohen and Perrault (1979); Grice (1957, 1968, 1969); Harris (2014); Neale (1992); Recanati (1987); Roberts (2012, 2017); Sbisá (2002); Schiffer (1972); Stalnaker (2014, 2017); Strawson (1964); Thomason (1990).

⁸Blume and Board (2013); Franke (2012); Harms (2004); Huttegger (2007); Millikan (1984, 1998, 2005); Skyrms (2010); Zollman (2011).

⁹E.g., Stalnaker (1999, 10–11; 1999, 10–11). Stalnaker sometimes casts his theory of assertion as a version of intentionalism. "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, 2017). However, in at least one place, he denies that all proposals to change the context need be intended to do so (1999, 87).

cillary aims. An intentionalist typically accomplishes this by saying that although an act may be intended to accomplish many things, the speaker *communicatively* intends to accomplish only a subset of those things, and only the effects that are communicatively intended can count as the act's essential aims.^{10,11}

The other crucial element of context-directed theories, as I am understanding them, is the theory of context that they rely on.¹² This theory can be spelled out by the following two principles.

(CP) CONTEXT PUBLICITY

Contexts are public: each of the participants in a conversation has equal epistemic access to the context.

(CR) CONTEXT REDUCIBILITY

Facts about context reduce to facts about the psychological states of the participants in a conversation.

The publicity of contexts is crucial given the roles that they are normally taken to play in conversation. For example, if it is common ground that the man in the corner is named 'Abe' (and there are no other salient Abes around), then a speaker who utters 'Abe' will be interpreted as referring to the man in the corner, and a speaker who wishes to refer to him might do so by uttering 'Abe'. This example illustrates two roles that common ground is thought to play. First, the information in the common ground is what hearers rely on as a supplement when linguistic convention does not fully specify what a speaker is saying. Second, because they know

¹⁰Grice (1957, 1969); Schiffer (1972). At a minimum, having a communicative intention requires intending both (i) to have some effect and (ii) for the addressee to recognize that one intends (i). Whether more than this is required has been the subject of elaborate debates; e.g.: Bach and Harnish (1979); Bennett (1976); Davis (1992, 2003); Grice (1957, 1969); Loar (1981); Neale (1992); Schiffer (1972); Sperber and Wilson (1995, 2015); Strawson (1964); Thomason (1990). I won't weigh in on this debate here.

¹¹A functionalist can say that although acts of a given kind typically have all sorts of effects, only some of these effects are functional, in that past instances of the effect caused instances of this type of act to be reproduced (1984; 2005; 2010).

¹²Contexts, thus conceived, go by various names: "common ground" and "context set" (Stalnaker, 1978, 2014), "discourse context" (Stalnaker, 1998), "conversational score" (Lewis, 1979), "files" (Heim, 1982, §3.1.4), "conversational record" (Thomason, 1990), "information structure" (Roberts, 2012), "information state" (Veltman, 1996), "conversational state" (Starr, ms, 2010), and so on. These different terms will have different connotations to those familiar with the literature, and some have used them to name different posits within the same theory (e.g. Camp 2017). I gloss over most of these subtleties here.

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 the common ground, will be interpreted in thus-and-such way. It's easy to see that common ground can play this role only if interlocutors all have equal and reliable epistemic access to it. When interlocutors disagree about what is common ground, the context is said to be defective (Stalnaker, 1978)—a state that, at least when relevant to the topic being discussed, will lead to miscommunication. So, insofar as we want to posit contexts at all, we want them to be public.¹³

I am skeptical that we need to hypostasize contexts to play these epistemic roles in interpretation, and I will say something about why in §3. However, my main point in this essay is not to deny that context plays these roles in the design and interpretation of utterances. Rather, I wish to deny that contexts play the conceptually distinct role of being the targets of communicative acts. It is crucial to understand that context could play the former role without playing the latter.

For Stalnaker, the publicity of contexts is intimately connected to their reducibility to interlocutors' private mental states.¹⁴ On Stalnaker's view, a proposition p is in the common ground if and only if the interlocutors commonly accept p for the purposes of the conversation: each accepts p , accepts that each of the others accepts p , accepts that each accepts that each accepts p , and so on. To accept p is to treat it as true for some purpose or other, either because one believes it or in some more provisional way.¹⁵ On this view, the publicity of common ground follows from its intersubjective reducibility: if an interlocutor doesn't accept p , or doesn't accept that others accept p (etc.), then p simply isn't in the common ground. It is impossible, by virtue of the nature of common ground, for one interlocutor to have more access to it than the others.

The idea that facts about context reduce to facts about interlocutors' propositional attitudes is attractive for several reasons. The strategy is parsimonious.

¹³For some variations on the view that context plays this closely related roles in allowing for the interpretation of communicative acts, see Asher and Lascarides (2003); Clark (1996); Clark and Carlson (1992); Clark and Marshall (1992); Clark et al. (1983); Ginzburg (2012); Lewis (1975, 1979); Portner (2007); Roberts (2002, 2003, 2005, 2012, 2017); Stalnaker (1978).

¹⁴The details of Stalnaker's theory of how common ground reduces to propositional attitudes have changed over time (1973; 1978; 1984; 1998; 2002; 2014). I rely on the most recent iteration of his view here (2014).

¹⁵Stalnaker thinks of acceptance as a "generic propositional attitude concept" that includes belief and various other more tentative belief-like attitudes, including supposition, and hypothesizing, that count as correct only when their contents are true (1984, 79–84). This makes common belief and common knowledge are special cases of common acceptance, for Stalnaker.

We presumably already have to believe that agents have private propositional attitudes; it would be nice not to have to also treat contexts as further theoretical primitives. More importantly, the psychological reducibility of contexts explains both the publicity of contexts and the role that contexts are commonly held to play in explanations of conversational dynamics. Much of the attraction of context-directed theories of communicative acts arises from their ability to explain why performing a communicative act at one stage in the conversation reliably influences how interlocutors act and interpret one another down the line. Because it is ultimately our private mental states—our beliefs, desires, intentions, and so on—that explain why we act in the ways we do, this kind of dynamic explanation presupposes an intimate connection between the context and private mental states.

To better understand this point, consider the most influential alternative to the idea that contexts are psychologically reducible, which is that contexts are “objective and normative” (DeVault and Stone, 2006; Lepore and Stone, 2015).¹⁶ On this view, the state of the context is not just whatever interlocutors take it to be; indeed, they could all be wrong about it. Instead, the state of the context at a given moment is determined by objective facts about the prior utterances that have been made in a discourse, together with facts about the linguistic conventions that govern those contributions. (This sort of view is most natural if paired with some version of dynamic semantics, though neither view is strictly necessary for the other.) As well as being “objective”, context is “normative” on this view because keeping track of it and acting in accordance with it is part of what participants should do, given the linguistic conventions in which they participate (DeVault and Stone, 2006).

The problem with this sort of view, as I see it, is that it creates the potential for situations in which facts about the state of context become entirely disconnected from explanations of how conversations actually play out. Imagine a situation in which two interlocutors, A and B, suddenly become massively mistaken about what has happened in the conversation up until then. Crucially, their mistake is a *folie à deux*: they are mistaken in exactly the same ways. So, for example, although it is an objective fact that the common ground contains p —say, because someone has previously asserted p —A and B now commonly accept $\neg p$ for the purposes of the conversation. Although a situation like this would presumably involve a considerable coincidence, and although we may want to consider it normatively defective in some way, nothing rules out the possibility of such a situation arising. How-

¹⁶Another alternative is that facts about context boil down to deeply normative facts about interlocutors’ commitments and entitlements (Brandom, 1994; Nickel, 2013). As far as I can see, this view is subject to the same kind of considerations that I’m about to raise for conventionalist theories of context.

ever, in a situation like this, the conversation would proceed as if $\neg p$ is common ground, not as if p is common ground. In other words: the body of propositions that matters for explaining how conversations will actually unfold is the one that is determined by interlocutors' mental states, not the objective and normative context posited by DeVault and Stone. More generally: facts about objective and normative contexts can explain how conversations unfold only insofar as they are tracked by psychologically reducible contexts. If our main goal in theorizing about context is to understand how conversations actually unfold, it is therefore psychologically reducible contexts (if any) that we should posit.¹⁷

I have been focusing on assertion and common ground, but analogous considerations extend to other components of contexts that have been posited by context-directed theorists. Roberts and Portner both argue that the components of context that they posit influence how interlocutors will design and interpret communicative acts. Roberts (2012) argues that communicative acts will be interpreted as relevant only if their contents bear certain logical relations to the QUD, for example. And Portner (2004; 2007) argues that deontic modals will normally be evaluated under the assumption that worlds are more closer to ideal if they realize more of the properties on the TDL. These roles are highly analogous to the pragmatic roles that Stalnaker assigns to common ground. In order to play them, facts about the QUD and TDL must be public. Once again, the best explanation of this is that these components of context are psychologically reducible in a way that gives all interlocutors equal epistemic access to them. Roberts and Portner seem to agree. In answering a question, Roberts says, "it is the common ground, not the speaker, that's 'informed,' and [in asking a question,] it is mutual-belief behavior, and not knowledge, that's sought" (Roberts, 2012, 6, fn.7). Likewise, Portner says that we can "think of the Common Ground and To-Do List as being the public, or interactional, counterparts of the individual agent's beliefs and desires" (Portner, 2004, 242). Of course, the question of precisely how facts about QUDs and TDLs reduce to facts about interlocutors' private mental states is an interesting and open one: neither Roberts, nor Portner, nor (as far as I know) anyone else has attempted to give a detailed reduction of the sort that Stalnaker offers for common ground. However, for present purposes I will assume that a necessary condition on the QUD or TDL being in a certain state is that the interlocutors commonly accept that it is in that state.¹⁸

¹⁷Of course, we might also have other theoretical goals. For example, we might want to explain why certain conversational trajectories are normatively defective. Objective and normative contexts might be useful for this purpose. For some other reasons for positing objective contexts, see Camp (2017); McGowan (2017).

¹⁸This formulation is somewhat problematic because it entails that language users must possess

2 Communication without Context Change

I turn now to my objections to context-directed theories. The central observation behind these objections is that it is possible to succeed in communicating a message to someone without making public either the content of the message or the fact that it has been sent, and without even having aimed at making one's message public in either of these ways. But, by (CP)+(CR), states of the context are constitutively public. And communicative acts are whatever a speaker has to do, and whatever an audience has to interpret, in order for communication to succeed. It follows that the constitutive aim of a communicative act can't be to change the context.

I begin with an example modeled on the coordinated-attack problem.¹⁹ Although this scenario is artificial, it has the virtue of being clear. In it, two generals are perched on hills straddling a valley in which their common enemy is encamped. If they attack at the same time, they will defeat the enemy, and this is their goal. But uncoordinated attacks will result in a defeat so severe that they must avoid that outcome at all costs, and so neither is willing to attack unless they have a high degree of certainty that they will do so in a coordinated way. The hitch is that the only channel by which they can communicate is unreliable: they can send messengers to each other through the valley, but there is always a high chance that a messenger will be captured. So, in this case, suppose that General A sends the message 'the attack will be at dawn' to General B, who receives it. Since he doesn't know that B got the message, A won't attack at dawn, and since B knows this about A, he won't attack at dawn either. B might try to improve their epistemic state by sending a messenger back to A with the message, 'I got your message. The attack will be at dawn.' This time, B doesn't know that A got the message, and so won't attack at dawn. And even if A does get the message, he knows that B doesn't know that he got it, and so also won't attack. This process can be iterated as much as we like, and the problem will never go away. Given plausible assumptions about their epistemic standards and risk aversion, it can be shown both that the generals won't attack unless they commonly know that they will attack at the same time, and that

the concept of a QUD or a TDL. So perhaps the following weaker constraint will suffice for my purposes: for any fact φ about the QUD or TDL, there exists a proposition p about the interlocutors' propositional attitudes such that the fact that p obtains together with the fact that p is commonly known by the interlocutors totally grounds φ .

¹⁹Some influential discussions of this problem include Akkoyunlu et al. (1975); Cohen and Yemini (1979); Fagin et al. (1995, 1999); Gray (1978); Halpern and Moses (1990); Moses (1986); Yemini and Cohen (1979). The closest precursor to my discussion of the coordinated-attack problem is Jankovic (2014), who uses it to object to theories of communicative acts that are similar to my targets (e.g., those of Schiffer 1972; Sperber and Wilson 1995), though for rather different ends.

they cannot ever achieve common knowledge of the time of their attack by means of this sort of unreliable channel. Indeed, in this case, it seems clear that the generals cannot come to commonly accept that they will attack at dawn, since they cannot reach a state in which they are both prepared to ignore, for the purposes of the conversation, the possibility that they won't attack at dawn. In other words, it is impossible for the generals to succeed in getting the proposition that they will attack at dawn into the common ground.

Does this show that the two generals can't communicate? I submit that it does not. The right way to describe the case would be to say that each message they send constitutes a communicative act, and that each time one of them receives a message, successful communication has taken place. Making their attack time common ground is a further, futile goal that they have—one that goes beyond communication and that they intend to accomplish *by means of* communication.

To see that this is the right description, suppose that, after sending and receiving several messages to B, A comes to realize that their attempt to coordinate is pointless, and sends one last message. A says that he has been reading some old papers by theoretical computer scientists and—alas!—they've proven that people in a situation like this will never be able to coordinate their attack, so they should just stay put. While he's writing, A adds, he has a question for B:

- (1) Some of my men have come down with cholera. What should we do?

General B receives this message, and responds, saying that it's too bad about the attack. As for the cholera:

- (2) You should tell your men to stop eating and drinking near the latrines.

Against all odds, A receives this last message and, as a result, stops his men from eating and drinking near the latrines. The cholera epidemic subsides.

Despite the elaborate setup, I see no reason to think that things couldn't play out in the way described. Moreover, there can be no doubt that, were it to occur, A would have asked a question that resulted in successful communication in uttering (1) and B would have performed an assertion (or perhaps a directive) that resulted in successful communication by uttering (2). How else, after all, are we to explain the fact that A gained a new piece of information about how to prevent cholera as a result of receiving A's message? This sort of reliable and precise information sharing is the exact phenomenon that a theory of communication must explain. But this suggests that communication was perfectly possible via their previous messages as well. What made it seem otherwise was that their broader goal

in communicating was to gain common knowledge for the purposes of coordinating their attack, and this broader goal was destined to fail. But, in general, one doesn't fail to communicate just by failing at achieving one's broader aims in communicating. From the fact that I don't get you to leave my apartment by saying 'I'm tired', it doesn't follow that I haven't communicated to you that I am tired, for example. What's happening in the coordinated-attack case, I submit, is the same phenomenon: A and B are communicating, but are unable to thereby achieve their broader, extra-communicative aims.

Moreover, as (1)–(2) show, communication for the purpose of making one's message common ground is a special case: it's possible to successfully communicate p even when one doesn't intend to make p common ground, either because one simply doesn't care to do so, or because one knows that it isn't possible to do so. In uttering (2), B knows (and knows that A knows) that he won't be able to get his message into the common ground. But B doesn't care about that, since getting his message into the common ground doesn't matter for either A's or B's broader purposes. The more important goal, for current purposes, is getting A to *believe* the content of his assertion, since having that belief will allow A to solve his cholera problem.

The coordinated-attack scenario is admittedly an artificial one for the purposes of studying human communication, and so it may be tempting to think that what I have just described is a paradox that affects only certain bizarre or hypothetical edge cases, and that the proper way of handling such cases is to ignore them or idealize away from them until we have a solid theory of the core cases of human communication. However, I don't think this is the right way to react. The reason is that many actual cases of successful communication have the same structural features. Call a situation *publicity averse* if it is one in which a speaker can't make the content of their communicative act public, and knows this. I think real people regularly find themselves publicity-averse situations but communicate anyway.

One macabre kind of publicity-averse situation arises whenever someone performs a communicative act that they intend to be understood only after they're dead. Consider the act of asserting something in a suicide note or directing one's executor to do something via a will, for example. In either case, the speaker clearly doesn't aim to make anything common ground between themselves and their addressee. After all: if things go according to plan, the speaker won't be around to engage in shared attitudes of any kind by the time their message is received. Their aim can only be to affect their addressee's private mental states. By doing so, people routinely succeed in communicating from beyond (or before) the grave.

A distressingly common kind of publicity-averse situation occurs whenever

one sends an email to a recipient whose inbox routinely overwhelms them for one reason or another.²⁰ I have several friends whose unread counts are larger than their salaries, and whose likelihood of actually opening and reading any given email they receive is surely as low as the chance of a courier making it through an enemy-infested valley. When I assert something in an email to one of these acquaintances, I have good reasons to doubt that they will read it, and so I can't hope to get the content of my assertion into the common ground. I communicate with them just fine when they do happen to open my emails, however, which is why I keep sending them.

These claims needn't rest entirely on our intuitions about when something is common ground, or when communication has succeeded. They are also supported by facts about how situations involving risky communicative acts often play out. Suppose that I am planning to sit in on the seminar of one of my inbox-overwhelmed friends. I send him an email that includes (3).

(3) I will be leaving your seminar early today to travel to a conference.

I then discover that the conference is cancelled. A short time later, I see my friend in person and I want to inform him of the cancellation. Here is a natural way for our conversation to go:

(4) ME: Did you read my email?

(5) FRIEND: Yes, I read it.

(6) ME: The conference is cancelled.

Although I might have simply led with (6), it is also natural for me to first begin with (4), particularly if I place a high value on avoiding misunderstanding. But leading with (4) makes no sense if the content of my email is already common ground: uttering (4) can be thought of as a way of repairing the common ground to get the content of my original message into it.²¹ However, the fact that merely confirming that my friend has received the email can suffice for this kind of repair

²⁰The point that email exchanges can be relevantly similar to the coordinated-attack scenario has been discussed in the context of the 'electronic mail game' (Lederman, 2017; Rubinstein, 1989). Jankovic (2014) considers the case in a way that parallels my argument here.

²¹Strategies for repairing the common ground of this kind are sometimes called 'grounding' or 'metacommunication' (Clark, 1996; Clark and Brennan, 1991; Clark and Schaefer, 1989; Ginzburg, 2012).

demonstrates that he has already understood my original message—i.e., that we've communicated. After all: I didn't tell him what was in my message, but merely checked that he got it; he has to know what its content was in order to put that content into the common ground, rendering my subsequent use of 'the conference' felicitous. We might also imagine that my friend has taken various actions on the basis of understanding my message. He might have moved topics relevant to me to the beginning of his lesson plan, for example. This would need to be explained by saying that we've already communicated.

The upshot, once again, is that successfully communicating a proposition does not require getting it into the common ground. And since I was aware of my friend's inbox problems before I sent my message, I knew that I couldn't get its content into the common ground just by sending it, and so I couldn't have rationally intended to get my message's content into the common ground either.²² And, once again, it may have been rational for me to send the message nonetheless. We can suppose that, when I sent my email, it was not all that important to me to get its content into the common ground; what mattered to me was getting my friend to believe that I had a good reason to leave his class early, so that he wouldn't be offended when I got up and left.

Publicity-averse scenarios like this are not unusual. Indeed, we should appreciate how ubiquitous they are. Any situation in which a speaker believes that the chance of their communicative act being understood is sufficiently low, but in which they perform it anyway, has the relevant qualities. This includes many situations involving communication via writing, but also situations involving oral communication in which the speaker isn't sure if the hearer is listening, or in which the speaker uses a language or vocabulary that they doubt the hearer will understand. I have been in many situations of this kind. When traveling in non-English-speaking countries, I sometimes tell a waiter that I am vegetarian, have serious doubts about whether they have understood me, and—out of some combination of timidity, politeness, and resignation—I go on doubting that I succeeded until my meal arrives. Sometimes I wind up with a customized, meat-free version of the dish I ordered, making it clear that I'd succeeded in communicating with the waiter after all. But this success can't have been constituted by making my preference for meat-free food common ground. Similarly, when giving a philosophy talk I sometimes use jargon that I doubt some of the audience members will understand. But from time

²²The principle that one can't rationally intend to ψ unless one believes that it is at least possible for one to ψ is a widely accepted principle governing the coherence of beliefs and intentions. For various versions of the principle, any of which would suffice for my purposes, see Bratman (1987); Broome (2013); Donnellan (1968); Grice (1971); Holton (2011); Neale (2004).

to time one of those audience members raises an issue during Q&A that reveals that they not only clearly understood what I was saying, but discovered a serious problem with it.

Of course, in many such situations—particularly when we want to continue the conversation—we immediately follow up on communicative acts like these immediately in order to repair the common ground. But this followup normally happens *after* the hearer has understood the content of the speaker’s message. In other words: all of these situations are temporarily publicity averse, if only for a short time. We can see this by focusing on publicity-averse situations in which followup is not possible, or merely delayed.

Suppose that I am boarding a train and my wife has come to the platform to see me off. After I’m in the train, I look out the window and see her. She has headphones in her ears and I believe that she’s already begun listening to music, and so can’t hear anything. Just as the train is pulling away, I remember something that I forgot to tell her, and, taking a chance, shout (7) before immediately losing sight of her, and without seeing her reaction.

(7) Please water my plant when you get home!

Suppose that my wife’s headphones happen to have been silent when I shouted. She understood my request and watered my plant when she got home. Again: we have to explain the fact that she does so by saying that she understood my request; we’ve communicated. However, I go on thinking that it is very unlikely that she understood me. Indeed, I believe that I failed. When I call her from my destination, one of the first things I say is (8), in an effort at repairing.

(8) I shouted about my plant from the train. Did you hear me?

This is a clear case of communication that succeeds in the absence of context update. But the only difference between this case and all other cases in which followup is needed is that this situation stays publicity averse for longer—long enough to allow behavioral evidence of the successful communication to emerge before followup can take place. Quite generally, I think, we should see the widespread existence of followup of this kind as evidence of successful communication in publicity-averse situations.

From these considerations, I conclude that successful communication requires producing certain changes in the addressee’s private mental state, but not changes to contexts, understood as public bodies of information that are grounded in interlocutors’ shared propositional attitudes. Although we often aim to change the

context when we communicate, this is not an essential feature of communication but an optional add-on that we aim for in some cases but not others, that we may fail at without failing to communicate, and that therefore isn't part of the essential aim of performing a communicative act.

3 Ways that Communicative Acts Can Succeed

The reader may now be thinking that I have been running together more than one sense in which a communicative act can turn out to be successful. It is true that I have been a bit sloppy about this so far. However, I don't think that my conclusions rest on any equivocation. In this section, I will try to demonstrate this by distinguishing between seven different ways that a communicative act can succeed, according to a Gricean, addressee-directed theory. Whereas addressee-directed theories predict that these should be real distinctions, context-directed theories elide several of them. I will argue that the distinctions thus elided are important ones.

In order to keep things simple and make the comparison between audience-directed and context-directed theories as straightforward as possible, I will focus on a genus of communicative acts that is unusually broad from the point of view of most Griceans. These are the communicative acts that one performs by producing an utterance with the communicative intention of getting one's addressee to *accept* some proposition. I follow Stalnaker (1984, 79–84) in taking acceptance to be a family of propositional attitudes that includes belief, supposition, hypothesis, and so on.²³ Assertion, as Grice thinks of it, is a species within this genus. This is a point that Stalnaker sometimes acknowledges by saying that assertion, as he defines it, is a broader notion than the one that is often discussed by philosophers (Stalnaker, 2017). So, let's call the genus of communicative act I am discussing 'assertion*'. I focus on the case of intending to produce acceptance rather than intending to produce belief because doing so allows for a more direct comparison with context-directed theories: on Stalnaker's account of assertion, the essential aim of asserting p is to produce a state of commonly accepting p among the participants in the conversation. Stalnaker's essential aim therefore asymmetrically

²³Some complications lurk here. For example, we may want to say that S can intend for A to accept p by intending for A to bear one of the various acceptance attitudes to p , in which case one needn't have the concept of acceptance to intend for someone to accept something. Since precisely parallel issues arise for Stalnaker's talk of intending or proposing to produce states of common acceptance, I will assume that nothing hangs on these issues.

entails the communicatively intended effect of assertion* as I have defined it. I will ignore directive acts and questions in order to avoid complicating the definitions below, though parallel considerations apply to them.

So, suppose we have a speaker *S* addressing an utterance *x* to an addressee *A*, and thereby asserting* *p*. Here, according to a Gricean, are seven ways in which such an assertion could succeed, along with corresponding success conditions.

- (S1) ILLOCUTIONARY SUCCESS
S utters *x* intending:
 - (i) for *A* to accept *p*; and
 - (ii) for *A* to believe that *S* uttered *x* intending for *A* to accept *p*.
- (S2) UPTAKE (COMMUNICATIVE SUCCESS)
A accepts that *S* uttered *x* intending for *A* to accept *p*.
- (S3) ACCEPTANCE (NARROW PERLOCUTIONARY SUCCESS)
A accepts *p*.
- (S4) WIDE PERLOCUTIONARY SUCCESS
A enters further mental states or takes further actions, partly a result of accepting *p*, thus fulfilling one or more of *S*'s goals in uttering *x*.
- (S5) COMMON UPTAKE
S and A commonly accept that *S* uttered *x* intending for *A* to accept *p*.
- (S6) COMMON ACCEPTANCE
S and A commonly accept *p*.
- (S7) PRESUPPOSITION LICENSING
Utterances presupposing *p* that follow *x* in the same discourse are felicitous.

Context-directed theorists typically focus on cases in which S1–S7 all occur, though cases in which all except for S6 and S7 occur are sometimes discussed in the context of rejected speech acts.²⁴

²⁴e.g. Thomason 1990; Stalnaker 1978; 1998). Building on work by Clark and Schaefer (1989) and others, Ginzburg (2012) offers the clearest discussion within the dynamic tradition of cases in which S5 and S6 fail to occur due to full-blown miscommunication. However, Ginzburg considers this phenomenon only as a kind of failure to communicate, and his focus is on understanding meta-communicative strategies for overcoming this failure. He models this phenomenon by taking each interlocutor to have a separate context (“Dialogue GameBoard” or DGB) that they seek to align but

Addressee-directed intentionalism predicts that each of these kinds of success is distinct. In what remains of this section, I will argue that this is a good prediction—i.e., that we really do need to distinguish all of these kinds of success in order to make sense of how humans use language. Moreover, as I began to argue in §2, we should think that the nature of communication is more intimately tied to S1–S3 than to subsequent forms of success.

First, notice that with the possible exception of S6 being sufficient for S7, none of these forms of success entails the subsequent forms. So, for example, one can succeed in performing a communicative act without anyone understanding it in any way: S1 does not invariably lead to any of S2–S7. After all: the thing that an addressee fails to understand in a case of misunderstanding is a communicative act.²⁵ It should also be clear that S2 does not invariably lead to S3 and that S3 does not invariably lead to S4: one can understand a communicative act without being convinced or persuaded by it, and one can be convinced or persuaded by it without drawing further conclusions or taking further action as a result. Finally, it should be clear, for analogous reasons, that S5 does not invariably lead to S6: even when it is commonly accepted that a communicative act has been performed, according to dynamic theorists, the act may still be rejected, thereby failing to have its essential effect.

The interesting questions, for my purposes, are about the relationships between S2–S4, on one hand, and S5–S7, on the other. My arguments in §2 were designed to show that S2–S4 can take place without either S5 or S6, that performing a communicative act is constituted by the aim of achieving S3 rather than S6, and that successfully communicating may result only from S2 but not S5–S6. In a publicity-averse situation, a speaker cannot hope to achieve COMMON UPTAKE or COMMON ACCEPTANCE. However, as long as their communicative act is understood (perhaps against the odds), they can be understood by, convince, and influence the actions of their addressee in complex ways. A theory of communication needs to explain this. Moreover, in each of the publicity-averse scenarios I've discussed, the

that sometimes come apart. He thus does not distinguish S1–S4 as distinctive kinds of success that may be valuable on their own, and so does not consider cases in which speakers either don't aim for or don't care about S5–S6 (or any other kind of common acceptance).

²⁵Austin held that many or all illocutionary acts require uptake for their success (1962, 115–6), and some others have defended this view (Hornsby and Langton, 1998; Langton, 1993). It seems to me that my disagreement with these authors is just a verbal one. What they call 'attempting to perform an illocutionary act' is what I call 'performing an illocutionary act', and what they call 'performing an illocutionary act' is what I call 'successfully communicating via an illocutionary act'. Although the view I am defending here is incompatible with Austin's theory of speech acts in many other ways, I don't think the issue of uptake is particularly deep.

speaker cares more about achieving S2–S4 than about S5–S6. Indeed, in some of the scenarios, S5–S6 aren't part of the speaker's goals at all.

On the other hand, there cannot be a case of S6 without a corresponding case of S3, since S6 entails S3: it is impossible for S and A to commonly accept *p* unless A accepts *p*. For analogous reasons, S5 cannot occur without S2. This shows that success of the kind that interests context-directed theorists is a special case of success of the kind that addressee-directed theorists have focused on.

Moreover, several of the more advanced of these forms of success build, in some interesting sense, on the antecedent kinds. S2 normally follows from S1, as do S3 from S2 and S6 from S5. Indeed, although it is possible to imagine scenarios in which a speaker causes their addressee to accept *p* by asserting* *p* without first getting the addressee to understand that the speaker has asserted* *p* (i.e., a case of S3 without S2), this would be a strange fluke rather than the sort of language use that deserves a theoretical explanation. An analogous point goes for speech acts that achieve COMMON ACCEPTANCE (S6) without COMMON UPTAKE (S5).

I have stipulated that WIDE PERLOCUTIONARY SUCCESS (S4) comes as a result of ACCEPTANCE (S3). In many cases, it makes sense to single out S4 in this way. For example, if General A clears his troops away from the latrines as a result of getting General B's message about the cause of cholera, he does so because he accepts (i.e. believes) the content of General B's message, not because he merely understands the message. Likewise in most other cases: a speaker normally attempts to prompt the addressee to act by convincing or persuading them of something.

However, we might want to distinguish kinds of perlocutionary success that follow from S2, S5, and S6 as well. In some cases, merely understanding a communicative act may be enough to prompt the addressee to act, and the resulting action may even have been intended by the speaker. If I tell a bouncer that he is a pathetic excuse for a man, this may be enough to provoke him to fight me, even if he doesn't accept what I say, and provoking him in this way may have been part of my drunken and misguided aim in speaking. Similarly, the two generals will actually attack their enemy only when they are satisfied that they have achieved S6—i.e., only when they commonly accept (i.e., know) that the attack is scheduled for dawn.

Finally, consider S7, which has been a central explanandum for dynamic models of conversation. Since context-directed theories of communicative are one component of dynamic models of conversation, we should expect that this is a form of success that they will be particularly interested in explaining. According to an influential principle first proposed by Stalnaker, it is felicitous to presuppose a proposition *p* only if *p* is entailed by the common ground (Stalnaker, 1973, 1974). This

principle admits of exceptions, since some presuppositions that aren't already common ground can be accommodated (Lewis, 1979; Roberts, 2015a), and this makes the status of Stalnaker's felicity requirement somewhat questionable. However, if it is correct to think of Stalnaker's rule as an exception-tolerant requirement, then we should expect that a significant part of the reason that speakers aim to get propositions into the common ground is so that they can later presuppose those propositions, or build on them in other ways. (I will say more about this in §4.)

However, I think that reflecting on the kinds of publicity-averse cases that I discussed in §2 shows that it can sometimes be felicitous to presuppose p even if p has not made it into the common ground. Indeed, in at least some circumstances, performing an assertion* whose content is p may license future presuppositions of p (S7) even if it does not succeed in achieving common uptake (S5) or common acceptance (S6) of p . Suppose I send the following email to an unreliable friend, thinking that it probably won't be read:

- (9) I have a brother who lives in Nova Scotia. Next week is my family's big family reunion in New York, and he's coming to town for it. [...] My brother really likes fancy pizza, so I'm going to take him to Di Fara's.

Both 'he' and 'my brother' are definite NPs, and it is widely assumed that it is felicitous to use a definite NP only in contexts that are set up in the right way. In other words: definite NPs are usually taken to have presupposition requirements. For example, it is sometimes held that an utterance of a sentence that includes 'my brother' (in at least some positions) presupposes that the speaker has a brother.²⁶ A standard way to render a use of 'my brother' felicitous is to first assert that one has a brother, as in (9). In a similar vein, many dynamic-semantic and dynamic-pragmatic theories of anaphora hold that it is felicitous to use a definite NP only when a corresponding *discourse referent* has been established in the context, and that facts about which discourse referents have been established in a context are grounded, at least in part, in facts about the propositions that the interlocutors commonly believe or accept.²⁷ One way to establish a discourse referent is to use an indefinite NP, so that uttering 'I have a brother' will in many cases allow subsequent utterances of 'he' and 'my brother' to refer to my brother.

At first glance, this process of adding to and building on the context might appear to explain what's going on in (9). But this can't be quite right. After all: (9)

²⁶See, e.g., Clark and Marshall (1992).

²⁷Heim (1982, §3.1.4); Roberts (2003, 2005).

is uttered in a publicity-averse situation, and so although S1–S4 are possible, S5–S6 are not. Still, there seems to be no infelicity in the speaker’s use of ‘he’ and ‘my brother’. So, although *p*’s being common ground may be sufficient to render it felicitous to presuppose *p*, it is not necessary.²⁸ And likewise, if a definite NP depends for its felicity on the establishment of a discourse referent by an antecedent indefinite NP, then establishing a discourse referent with an indefinite NP does not require getting any proposition into the common ground (including the proposition that the indefinite NP in question was uttered). We can at least sometimes succeed at licensing a presupposition in a way that bypasses the context—in a way, that is, that does not involve COMMON UPTAKE (S5) or COMMON ACCEPTANCE (S6).

I wish to illustrate this point with one more example, lest the reader be tempted to think that it applies only in certain special cases involving writing. Professor X is teaching a logic class in which she first introduces the axiom of extension as follows.

- (10) Here is an axiom: two sets are equal if and only if they have the same members.

Prof. X goes on to make various remarks about the axiom. At various later points in the lecture, she refers to it as ‘the axiom’ or as ‘it’. Throughout the lecture, Prof. X notices that Chad, a student sitting in the front row, appears to be asleep: his eyes are closed, he is drooling a bit, and his head is tilted back and to the side. Prof. X concludes that Chad wasn’t awake when she introduced the axiom and that he therefore doesn’t accept the axiom nor the proposition that she expressed it. Therefore, neither the axiom nor the proposition that Prof. X asserted* it is common ground. When the Q&A comes, however, Chad surprises professor X by opening his eyes and asking the following astute question:

- (11) Am I right to think that it follows from the axiom that there is only one empty set?

Here, Chad felicitously uses a definite NP to anaphorically refer to a proposition introduced earlier in the discourse, despite the fact that neither the antecedent proposition itself nor the proposition that it had been expressed is common ground.

So, presupposition licensing does not require changing the common ground. This may go some way toward explaining why presuppositions that aren’t common ground can often be felicitously accommodated. It may be asked: what *is* necessary for presupposition licensing if not common acceptance or even common uptake?

²⁸The same goes for various other kinds of presupposition. There’s no problem about using ‘still’, ‘again’, it-clefts, and so on in a publicity-averse situation, for example.

I won't attempt to give a full answer to this question, but here is an initial thought. In order to use a definite NP felicitously, the only thing that seems essential is that (a) the presuppositions it triggers don't directly conflict with anything that the addressee accepts (or anything they accept about what the speaker accepts, etc.) in a way that the addressee is unwilling to revise without comment, and that (b) the addressee is capable of identifying the definite's referent, by whatever means are available to them. These two conditions follow from most common-ground-centric theories of the norms governing presupposition, but they are strictly weaker, since they don't give common ground any essential role to play.

This section has shown several important things. First, a theory of communication needs to traffic in distinctions that are available to addressee-directed theories, but that are elided by context-directed theories and the dynamic models of conversation of which they are one component. Second, context plays a less essential role in communication than is commonly thought. We can successfully communicate with people (S₂), convince them of things (S₃), get them to do things (S₄), and license downstream presuppositions and anaphora (S₇) without doing so by means of changing the context (S₅–S₆). However, the kinds of success involved in changing the common ground always requires some non-common-ground-involving forms of success (S₅ entails S₂ and S₆ entails S₃). Addressee-directed theories of communication are thus more general and better allow us to understand the factors that are essential to communication, as opposed to those that are present only in certain special cases.

4 Idealization and Conversational Dynamics

The reader may be tempted to respond to my arguments so far by protesting that context-directed models of communicative acts are inseparable from dynamic models of conversation, and that the latter are *idealized* models. Few proponents of these models would claim that they can explain every nuance or case of communication. Rather, they are meant to capture core cases, or cases that are particularly interesting for certain theoretical purposes, or cases that are distorted in merely harmless ways given our explanatory purposes. And, this line of thought goes, it is fatuous to object to an idealized model by merely pointing out that it is not realistic in some way: that's what idealization is all about.

I think that this line of thought does articulate the ambitions of dynamic models of conversation. However, I think it is worth reflecting on the nature, purpose, and limitations of the idealization involved. Science deals in different kinds of ide-

alization, and the legitimacy of an idealization depends on the explanatory goals to which it is put. What is the purpose of the idealization involved in context-directed theories of communication? And is the form of idealization involved in these theories a legitimate and useful way of pursuing this purpose? To answer these questions, it will help to consider some of the work that philosophers of science have done to clarify the nature and purposes of idealization.

One possibility is that dynamic models of communication involve what Weisberg (2007; 2013) calls “minimalist idealization”—the practice of eliminating all factors that fail to “make a difference” to the target phenomenon being studied. Drawing on work by Strevens (2011), Weisberg illustrates the idea of a minimalist model by discussing its use in Boyle’s law.

In explaining Boyle’s law, for example, theorists often introduce the assumption that gas molecules do not collide with each other. This assumption is false; collisions do occur in low-pressure gases. However, low-pressure gases behave as if there were no collisions. This means that collisions make no difference to the phenomenon and are not included in the canonical explanation. Theorists’ explicit introduction of the no-collision assumption is a way of asserting that collisions are actually irrelevant and make no difference. (Weisberg, 2007, 643)

Of course, the idea that some causal factors don’t make a difference or are irrelevant to the target phenomenon presupposes that we have a clear grasp of what the target phenomenon is. And, in general, the legitimacy of a minimalist idealization will be relative to the choice of target phenomenon. The kind of idealization involved in Boyle’s law might be appropriate when what we’re trying to understand is the macroscopic behavior of gases, but it would be inappropriate if our goal is to understand their microscopic behavior, for example.

If dynamic models of conversation are to be understood as partaking in minimalist idealization, and if what I’ve said in §§2–3 is correct, then the morals for dynamic models of conversation are relatively clear. If our goal is to understand the nature of communication or closely related notions such as linguistic meaning, then dynamic models of conversation idealize in illegitimate ways. After all: they don’t simply elide non-difference-making phenomena. Rather, they elide phenomena that are central to the nature of communication, along with plenty of real-world cases of communication. They thus distort the target phenomenon not in an irrelevant way, but in a highly relevant and misleading way.

The point can perhaps best be seen by focusing on dynamic-semantic theories, which identify the meanings of clauses with context-change potentials—mappings

between states of the context.²⁹ The way in which these theories purport to explain the role of linguistic meaning in communication is intimately linked to the idea that communication is a matter of changing the context. Uttering a sentence is a good way to perform a communicative act because performing a communicative act is a matter of changing the context in a certain way, and sentence meanings are instructions for doing just that. But, as I've shown, communication needn't aim at or result in context change, and linguistic meaning presumably plays a crucial role in *this* communication as well. Dynamic semanticists thus define linguistic meaning—a commodity that we rely on in the messy real world—in terms of features that are peculiar to their idealized models.

Perhaps dynamic models of conversation involve a different kind of idealization then. Another possibility is that they are *special cases* models—models that “directly describe one class of cases, which are simple and tractable, and use these as the basis for a more indirect understanding of the others. Understanding is achieved via similarity relations between the simple cases we have picked apart in detail, and the cloud of more complicated ones.” (Godfrey-Smith, 2009, 4).³⁰

Can dynamic models of conversation be understood as descriptions of particularly simple cases of communication that can then be used to understand more complex cases indirectly, by means of their similarity to the simple cases? I don't think so. Given what I have argued in §§2–3, there is no clear sense in which communication that involves context change is simpler than communication that does not involve context change. Indeed, cases of communication involving context change involves extra steps, over and above the simpler form of communication that is also present in non-context-involving cases. Moreover, it seems unlikely that mere similarity relations between the two kinds of cases will be enough to allow us to understand the non-context-involving cases in terms of the context-involving ones. Rather, what's needed to understand communication in publicity-averse cases is a different, more detailed theory of communication itself.

A more promising possibility is that context-directed models of communication involve what McMullin (1985) and Weisberg (2007; 2013) call *Galilean idealization*—“the practice of introducing distortions into theories with the goal of simplifying theories in order to make them computationally tractable” (Weisberg, 2007, 640). This kind of idealization takes its name from Galileo because he used

²⁹E.g. Beaver (2001); Groenendijk and Stokhof (1989, 1991); Heim (1982, 1983a,b); Kamp (1981); Kamp and Reyle (1993); Murray (2014); Murray and Starr (2017); Starr (ms, 2010); Veltman (1996).

³⁰My use of the term ‘special cases idealization’ is borrowed from unpublished work by Santana (MS), who uses Godfrey-Smith as a key example.

it liberally, for example, when ignoring the effects of air resistance in order to develop a model of projectile motion. “Galilean idealization”, Weisberg tells us, “is justified pragmatically” (2007, 641): although the model that results from such an idealization is a distorted one that sometimes makes inaccurate predictions, it is better than nothing, which is what we would have if we attempted to account for all of the factors that influence a phenomenon at once, without knowing how to handle some of them.

It makes some sense to think of context-directed models of communication as being idealized in this way. The reason is that these models tend to idealize away from aspects of linguistic communication that rely on non-monotonic reasoning.³¹ The processes by which each of the kinds of success discussed in the previous section give rise to subsequent forms of success all involve reasoning of this kind. It is a well-worn point that interpreting even the most direct and literal communicative acts—i.e., getting from S_1 to S_2 —involves recognizing the speaker’s intentions, for example.³² And intention recognition, being a kind of mindreading, is a paradigmatic example of non-monotonic reasoning. Likewise, the processes by which we decide whether to accept or reject a communicative act that we’ve understood ($S_2 \rightarrow S_3$), the process of deciding what else to think or do on the basis of new information ($S_3 \rightarrow S_4$), and the process of reasoning about how well others’ thoughts match one’s own ($S_3 \rightarrow S_5$; $S_4 \rightarrow S_6$) are all paradigmatically non-monotonic forms of reasoning.

All of this is relevant for current purposes because the computational intractability of non-monotonic reasoning is the best known obstacle to cognitive-scientific models of the human mind. This is the point of what Fodor calls his First Law of the Nonexistence of Cognitive Science: “the more global (e.g. the more isotropic) a cognitive process is, the less anybody understands it” (Fodor, 1983, 107). So, given that we don’t have realistic or predictive computational models of the processes by which the different forms of success are linked, perhaps it makes sense to elide the various steps in the processes that link the intermediate stages between S_1 and S_6 ? This is what most dynamic models of conversation do. By ignoring all the messy parts of communication that we don’t know how to formally model, we gain the ability to treat conversation as a rule-governed and deterministic process.

³¹A.k.a. ‘non-demonstrative inference’ (Sperber and Wilson, 1995), ‘inference to the best explanation’ (Harman, 1965), ‘central cognition’ (Fodor, 1983), or ‘abductive reasoning’.

³²See, e.g., Carston (2002); Heim (2008); King (2013, 2014); Neale (2004, 2007); Sperber and Wilson (1995). Even if we take the semantic values of context-sensitive expressions to be fixed by something other than the speaker’s intentions, the processes by which hearers identify those values will still presumably involve non-monotonic reasoning of some kind.

This is arguably the biggest strength of dynamic models of conversation. They offer us pleasingly concise and easily formalized models of how discourses evolve over time. A discourse can be thought of as a series of moves in a formally precise game, each of which is characterized by the way it changes the context, and in which each change to the context influences future moves in predictable, rule-governed ways. A communicative act, on this view, can be modeled as a mapping between states of the context (Gazdar, 1981; Lewis, 1979). Thinking of discourses in this way—as the product of what Rothschild and Yalcin (2016; 2017) call “conversation systems”—allows us to study them as formal objects, to prove theorems about them, to program computers to engage in them, and to do all sorts of other fun stuff that less idealized theories of conversation don’t allow for. According to Yalcin, this is the benefit we get from studying conversation at a very high level of abstraction (2017; see also Rothschild and Yalcin 2016; 2017).

This may be a worthwhile pursuit, but we mustn’t overestimate the explanatory benefits that an idealized model of this kind can offer. We mustn’t let the model’s cleanliness cause us to forget about the unwieldy mess that we’ve temporarily swept under the rug. After all, we’ll eventually need to lift our Galilean idealization—to de-idealize the model. Like a model of projectile motion that elides air resistance, a model of conversation that elides the distinctions between S₁–S₆ can only ever offer us a partial understanding and imperfect predictions concerning its target phenomenon. Eventually, to get a full picture of how conversations work, we will need to find ways to build those distinctions back in.

In fact, this analogy offers a view of dynamic models that is in some ways too optimistic. After all: as I argued in §2, there are perfectly usual and explanation-worthy cases of linguistic communication that don’t involve context change at all. Context-directed theories simply can’t make sense of these cases. And so, dynamic models of conversation not only idealize away from a variety of messy steps that are normally involved in conversation; they idealize away from a range of cases of communication as well. These two forms of idealization are connected. What is shared by the cases that fit dynamic models and cases that don’t fit dynamic models are precisely those forms of success that dynamic models elide. If, as I have argued, the essence of communicative acts and successful communication is a matter of aiming at and achieving these forms of success (S₃ and S₂), then it makes sense to say that dynamic models of conversation idealize away from communication itself, in order to focus on abstract properties of certain special kinds of conversation.

Yalcin (2017) may be happy to grant this. He argues, for example, that the notion of speech-act force that is used within dynamic models should be considered distinct from the traditional notion of illocutionary force, and native to a higher

level of abstraction at which, for example, the semantics–pragmatics distinction is ignored. But I worry that this level of abstraction is so precipitous as to deny us a good view of the action. Assuming that the goal of semantic and pragmatic theorizing is ultimately to explain how people do things with language—and in particular, how they communicate—a level of abstraction from which it is impossible to distinguish communication from non-communication would seem somewhat dubious in its explanatory scope. At the very least, we should demand a clear articulation of what its explanatory ambitions are.

Another way to see the same point is to recognize that the legitimacy of Galilean idealization, like minimalist idealization, is always relative to our explanatory goals. Just as a model of projectile motion that idealizes away air resistance would be an inappropriate tool for studying air resistance, dynamic models are inappropriate tools for studying the nature of communicative acts and communication. This is not to say that they can't be useful for some other purposes. But we must recognize their explanatory limitations and get clear on what these purposes could be.

5 Conclusion

Contrary to recent trends, communicative acts should not be understood as attempts to pool information with one's peers, or as attempts to coordinate with them in some way. On the contrary: communicative acts needn't aim at affecting the conversation's context in any way, and communication can be successful without changing the context. I have argued that we would do better to understand communicative acts in the way that Grice understood them, as attempts to change addressees' private states of mind, in part by revealing one's intention to do so. Even when communication does result in context change, this should be understood as an extra-communicative accomplishment, much as offending someone, seducing them, or provoking them to violence are extra-communicative accomplishments.

References

- Akkoyunlu, E. A., Ekanadham, K., and Huber, R. V. (1975). Some constraints and tradeoffs in the design of network communications. *ACM SIGOPS Operating Systems Review*, 9(5).
- Alston, W. P. (2000). *Illocutionary Acts and Sentence Meaning*. Cornell University Press.
- Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Cambridge University Press, Cambridge, UK.
- Austin, J. L. (1962). *How to do Things with Words*. The Clarendon Press, Oxford.
- Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.
- Bar-On, D. (2013). Origins of meaning: Must we ‘go gricean’? *Mind and Language*, 28(3):342–375.
- Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI, Stanford, CA.
- Bennett, J. (1976). *Linguistic Behaviour*. Cambridge University Press, Cambridge, UK.
- Blume, A. and Board, O. (2013). Language barriers. *Econometrica*, 81(2):781–812.
- Brandom, R. (1983). Asserting. *Noûs*, 17(4):637–650.
- Brandom, R. (1994). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.
- Brandom, R. (2000). *Articulating Reasons*. Harvard University Press.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.
- Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Camp, E. (2017). Slurs as dual-act expressions. In Sosa, D., editor, *Bad Words*. Oxford University Press, Oxford, UK.
- Carlson, L. (1982). *Dialogue Games: An Approach to Discourse Analysis*. Number 17 in Synthese Language Library. D. Reidel, Dordrecht, The Netherlands.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Blackwell.
- Clark, H. H. (1996). *Using Language*. Cambridge University Press, Cambridge, UK.
- Clark, H. H. and Brennan, S. E. (1991). Grounding in communication. In Resnick, L., Levine, B. M., and Teasley, S. D., editors, *Perspectives on Socially Shared Cognition*, pages 127–149. American Psychological Association.
- Clark, H. H. and Carlson, T. B. (1992). Context for comprehension. In Clark, H. H., editor, *Arenas of Language Use*, chapter 2, pages 60–77. University of Chicago Press and Center for the Study of Language and Information.
- Clark, H. H. and Marshall, C. R. (1992). Definite reference and mutual knowledge. In Clark, H. H., editor, *Arenas of Language Use*, chapter 1, pages 9–59. University of Chicago Press and Center for the Study of Language and Information.
- Clark, H. H. and Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, 13(2):259–294.
- Clark, H. H., Schreuder, R., and Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior*, 22:245–258.
- Cohen, D. and Yemini, Y. (1979). Protocols for dating coordination. In *Proceedings of the Fourth Berkeley Conference on Distributed Data Management and Computer Networks*, pages 179–188. Lawrence Berkeley Laboratory, Berkeley, CA.
- 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 Pinõn, C., editor, *Empirical Issues in Syntax and Semantics 9: Papers From the Colloque de Syntaxe et Sémantique à Paris 2011*, pages 37–58.
- Davis, W. (1992). Speaker meaning. *Linguistics and Philosophy*, 15(3):223–253.
- Davis, W. (2003). *Meaning, Expression, and Thought*. Cambridge University Press.
- DeVault, D. and Stone, M. (2006). Scorekeeping in an uncertain language game. In Schlangen, D. and Fernández, 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.
- Donnellan, K. S. (1968). Putting humpty dumpty together again. *Philosophical Review*, 77(2):203–215.
- Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning about Knowledge*. MIT Press, Cambridge, MA.
- Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1999). Common knowledge revisited. *Annals of Pure and Applied Logic*, 96(1):89–105.
- 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. (ms). A modest proposal for the meaning of imperatives. In Arregui, A., Rivero, M., and Salanova, A. P., editors, *Modality Across Semantic Categories*. Oxford University Press.
- Fodor, J. (1983). *The Modularity of Mind*. MIT Press, Cambridge, Mass.
- Franke, M. (2012). On assertoric and directive signals and the evolution of dynamic meaning. *International Review of Pragmatics*, 4:232–260.
- 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.
- Godfrey-Smith, P. (2009). *Darwinian Populations and Natural Selection*. Oxford University Press, Oxford and New York.
- Gray, J. (1978). Notes on data base operating systems. In Bayer, R., Graham, R. M., and Seemüller, G., editors, *Lecture Notes in Computer Science 60: Operating Systems*, pages 394–481. Springer-Verlag.
- Green, M. S. (2007). *Self-Expression*. Oxford University Press.
- 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.
- Grice, H. P. (1971). Intention and uncertainty. *Proceedings of the British Academy*, 57:263–279.
- Groenendijk, J. and Stokhof, M. (1989). Dynamic montague grammar. In Kálmán, L. and Pólos, L., editors, *Logic and Language*. Akadémiai, Budapest.
- Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.
- Gunlogson, C. (2001). *True to Form: Rising and Falling Declaratives as Questions in English*. Routledge.
- Halpern, J. Y. and Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587.
- Hamblin, C. L. (1971). Mathematical models of dialogue. *Theoria*, 37:130–155.
- Harman, G. (1965). The inference to the best explanation. *The Philosophical Review*, 74(1):88–95.
- Harms, W. F. (2004). *Information and Meaning in Evolutionary Processes*. Cambridge University Press, Cambridge, UK.
- Harris, D. W. (2014). *Speech Act Theoretic Semantics*. PhD Dissertation, City University of New York Graduate Center.
- Harris, D. W., Fogal, D., and Moss, M. (2017). 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. (1983a). 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.
- Heim, I. (1983b). On the projection problem for presuppositions. In *Second West Coast Conference on Formal Linguistics (WCCFL 2)*, pages 114–125.
- Heim, I. (2008). Features on bound pronouns. In Harbour, D., Adger, D., and Bejar, S., editors, *Phi-Theory: Phi-Features across Modules and Interfaces*. Oxford University Press, Oxford, UK.
- Holton, R. (2011). *Willing, Wanting, Waiting*. Oxford University Press.
- Hornsby, J. and Langton, R. (1998). Free speech and illocution. *Legal Theory*, 4(1):21–37.
- Huttegger, S. M. (2007). Evolutionary explanations of indicatives and imperatives. *Erkenntnis*, 66:409–436.
- Jankovic, M. (2014). Communication and shared information. *Philosophical Studies*, 169(3):489–508.
- 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.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.
- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27.
- King, J. C. (2014). Speaker intentions in context. *Noûs*, 48(2):219–237.
- Kukla, R. (2014). Performative force, convention, and discursive injustice. *Hypatia*, 29(2).
- Kukla, R. and Lance, M. (2009). 'Yo!' and 'Lo!': *The Pragmatic Topography of the Space of Reasons*. Harvard University Press.
- Langton, R. (1993). Speech acts and unspeakable acts. *Philosophy and Public Affairs*, 22:305–330.
- Lederman, H. (2017). Two paradoxes of common knowledge: Coordinated attack and electronic mail. *Noûs*, forthcoming.
- Lepore, E. and Stone, M. (2015). *Imagination and Convention*. Oxford University Press, Oxford.
- Lewis, D. K. (1975). A problem about permission. In et al., E. S., editor, *Essays in Honour of Jakko Hintikka*. D. Reidel, Dordrecht, The Netherlands.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Loar, B. (1981). *Mind and Meaning*. Cambridge University Press, Cambridge, UK.
- MacFarlane, J. (2011). What is assertion? In Brown, J. and Cappelen, H., editors, *Assertion*, pages 79–96. Oxford University Press.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- McGowan, M. K. (2017). On covert exercitives: Speech and the social world. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science*, 16(3):247–273.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press.

- Millikan, R. G. (1998). Proper function and convention in speech acts. In Hahn, L. E., editor, *The Philosophy of Peter F. Strawson*, The Library of Living Philosophers, pages 25–43. Open Court, LaSalle, Illinois.
- Millikan, R. G. (2005). *Language: A Biological Model*. Oxford University Press.
- Moses, Y. (1986). *Knowledge in a Distributed Environment*. PhD thesis, Stanford, Stanford, CA.
- Murray, S. E. (2014). Varieties of update. *Semantics and Pragmatics*, 7(2):1–53.
- Murray, S. E. and Starr, W. B. (2017). Force and conversational states. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Murray, S. E. and Starr, W. B. (MS). The structure of communicative acts. *Unpublished Manuscript, Online at <https://www.dropbox.com/s/ctf8stqvpl1rdqs/The-Structure-of-Communicative-Acts.pdf?dl=1>*.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15:509–559.
- Neale, S. (2004). This, that, and the other. In Bezuidenhout, A. and Reimer, M., editors, *Descriptions and Beyond*, pages 68–182. Oxford University Press.
- Neale, S. (2007). Heavy hands, magic, and scene-reading traps. *European Journal of Analytic Philosophy*, 3(2):77–132.
- Nickel, B. (2013). Dynamics, brandom-style. *Philosophical Studies*, 162(2):333–354.
- Pagin, P. (2011). Information and assertoric force. In Brown, J. and Cappelen, H., editors, *Assertion*, pages 97–136. Oxford University Press.
- 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.
- Portner, P. (2017). Commitment to priorities. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford, UK.
- Recanati, F. (1987). *Meaning and Force: The Pragmatics of Performative Utterances*. Cambridge University Press, Cambridge, UK.
- Roberts, C. (2002). Demonstratives as definites. In van Deemter, K. and Kibble, R., editors, *Information Sharing*, pages 1–48. CSLI.
- Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26:287–350.
- Roberts, C. (2004). Context in dynamic interpretation. In Horn, L. and Ward, G., editors, *The Handbook of Pragmatics*, pages 197–220. Blackwell.
- Roberts, C. (2005). Pronouns as definites. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 503–543. Oxford University Press, Oxford, UK.
- Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.
- Roberts, C. (2015a). Accommodation in a language game. In Loewer, B. and Schaffer, J., editors, *A Companion to David Lewis*, pages 345–366. Wiley-Blackwell, Oxford, UK.
- Roberts, C. (2015b). 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. (2017). Speech acts in discourse context. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*, Oxford, UK. Oxford University Press.
- Rosenthal, D. M. (1986). Intentionality. *Midwest Studies in Philosophy*, 10:151–184.
- Rothschild, D. and Yalcin, S. (2016). Three notions of dynamicness in language. *Three Notions of Dynamicness in Language*, 39(4):333–355.
- Rothschild, D. and Yalcin, S. (2017). On the dynamics of conversation. *Noûs*, 51(1):24–48.
- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *The American Economic Review*, 79(3):385–391.
- Santana, C. (MS). Intention-based accounts of meaning are idealized scientific models. *Unpublished Manuscript*.
- Sbisá, M. (2002). Speech acts in context. *Language and Communication*, 22(4):421–436.
- Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.
- Searle, J. (1965). What is a speech act? In Black, M., editor, *Philosophy in America*, pages 221–239. Allen and Unwin, London.
- Searle, J. (1969). *Speech Acts*. Cambridge University Press, London.
- Searle, J. and Vanderveken, D. (1985). *Foundations of Illocutionary Logic*. Cambridge University Press.
- Skyrms, B. (2010). *Signals*. Oxford University Press.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford, 2 edition.
- Sperber, D. and Wilson, D. (2015). Beyond speaker meaning. *Croatian Journal of Philosophy*, XV(44):117–150.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2:447–457.
- Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. and Unger, P., editors, *Semantics and Philosophy*. New York 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*, 25:701–721.
- Stalnaker, R. (2014). *Context. Context and Content*. Oxford University Press, Oxford.
- Stalnaker, R. (2017). 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.
- Strawson, P. F. (1964). Intention and convention in speech acts. *The Philosophical Review*, 73(4):439–460.
- Strevens, M. (2011). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.
- 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.
- Weisberg, M. (2007). Three kinds of idealization. *Journal of Philosophy*, 104(12):639–659.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press, Oxford and New York.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105:489–523.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.
- Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160.
- Yalcin, S. (2017). Expressivism by force. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press.
- Yemini, Y. and Cohen, D. (1979). Some issues in distributed process communication. In *Proceedings of the 1st International Conference on Distributed Computing Systems*, pages 199–203.
- Zollman, K. J. S. (2011). Separating directives and assertions using simple signaling games. *The Journal of Philosophy*, 108(3):158–169.