



# We talk to people, not contexts

Daniel W. Harris<sup>1</sup> 

Published online: 18 September 2019  
© Springer Nature B.V. 2019

**Abstract** According to a popular family of theories, assertions and other communicative acts should be understood as attempts to change the context of a conversation. Contexts, on this view, are publicly shared bodies of information that evolve over the course of a conversation and that play a range of semantic and pragmatic roles. I argue that this view is mistaken: performing a communicative act requires aiming to change the mind of one’s addressee, but not necessarily the context. Although changing the context may sometimes be among a speaker’s aims, this should be seen as an extra-communicative aim, rather than one that is necessary for the performance of a communicative act. Along the way, I also argue that contexts needn’t play a role in linking anaphora to their antecedents. On the view that I defend, theories that take publicly shared contexts to play an essential role in the nature of communicative acts or anaphoric dependence conflate an artifact introduced by idealized models of conversation with a feature of the phenomenon being modeled.

**Keywords** Speech acts · Communication · Context · Pragmatics · Stalnaker · Grice · Anaphora · Assertion

---

✉ Daniel W. Harris  
danielwharris@gmail.com

<sup>1</sup> Hunter College, CUNY, 695 Park Ave., New York, NY 10065, USA

## 1 Introduction

A communicative act is the kind of speech act that a speaker has to perform and their addressee has to correctly interpret in order for communication to happen.<sup>1</sup> Paradigm cases are assertions, questions, and requests. My aim here is to compare two kinds of theory about the nature of communicative acts, one trendy and the other a bit old-fashioned, and then defend the old-fashioned option.

I will call theories of the trendy kind ‘context-directed theories’. According to them, to perform a communicative act is to do something whose essential aim is to change the conversation’s context. For the purposes of these theories, which often come packaged with dynamic models of conversation, the context of a conversation is specific kind of posit—a body of information that is publicly shared by the participants in a conversation.<sup>2</sup> An influential context-directed theory is Stalnaker’s (1978, 2014) account of assertion, according to which asserting  $p$  is an act whose essential aim is to add  $p$  to the common ground. The common ground, according to Stalnaker, is the set of propositions that the interlocutors commonly accept for the purposes of the conversation. Other kinds of communicative acts have been understood in analogous ways, except that they are essentially aimed at changing different components of the context.

The old-fashioned alternatives are what I’ll call ‘addressee-directed theories’. According to these theories, the essential aim of a communicative act is to change the addressee’s state of mind in some way. For example, according to Grice (1957, 1968, 1969), to assert  $p$  is to do something with the intention of getting one’s addressee to believe  $p$ , in part by getting them to recognize that this is what one intends.<sup>3</sup> Communicative acts of other kinds are understood analogously, but are aimed at changing mental states other than beliefs.<sup>4</sup>

Addressee-directed theories are compatible with the existence of contexts, and with the idea that we sometimes aim to influence contexts by performing communicative acts. However, if an addressee-directed theory is correct, then changing the context is at best a secondary and inessential aim of communicative

---

<sup>1</sup> Following Bach and Harnish (1979), I distinguish communicative acts from conventional acts, such as performing a marriage ceremony or testifying in court, which can be performed only against the background of social conventions that define the conditions of their performance. Although some speech-act theorists would deny this distinction (e.g. Austin 1962), it is uncontroversial according to the views under discussion here.

<sup>2</sup> It is important to distinguish between a dynamic model of conversation and any particular proposal about how this model maps onto and illuminates human language use. This essay criticizes dynamic models only insofar as they are interpreted in terms of context-directed accounts of communicative acts and public, intersubjective theories of context. This is a widespread interpretation but not the only possible one.

<sup>3</sup> Many elaborations on this definition have been explored. Aside from those that effectively turn Grice’s theory into a context-directed theory, I will bracket most of these elaborations here.

<sup>4</sup> Grice used the terms ‘utterer’s occasion meaning’ and ‘indicative-type utterance’ rather than ‘communicative act’ and ‘assertion’, respectively. I use the latter terms because they make it easier to compare Grice’s view with other theories of communicative acts, and because they make it easier to discuss non-assertoric communicative acts. In this respect, I follow Bach and Harnish (1979) and various others.

acts. To change the context is something over and above the kind of success at which a communicative act must be aimed in order to qualify as a communicative act.

These two views have a lot in common. Both take the nature of a communicative act to be a matter of its essential 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 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 (Harris et al. 2018). I bracket these other views here.

My plan is to defend addressee-directed views by showing that they make better sense of many cases of communication (Sect. 3), by distinguishing a variety of ways in which a communicative act can succeed (Sect. 4), and by arguing that context-directed theories confuse the features of idealized models with the features of their target subject matter (Sect. 5). First, in Sect. 2, I will clarify the debate.

## 2 Context-directed theories of communicative acts

The central thesis of context-directed theories can be stated as follows.<sup>5</sup>

### (CD) CONTEXT DIRECTEDNESS

The essential aim of a communicative act is to change the context in a particular way.

A well known theory of this kind is Stalnaker's theory of assertion, which I sketched in the last section. Others have posited additional components of contexts for other kinds of communicative acts to manipulate. Roberts (2012, 2004, 2018) argues that the context includes a *question-under-discussion stack*, that the question at the top of the stack 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.<sup>6</sup> Similarly, Portner (2004) 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 directive act is to add a new commitment to the addressee's section of the TDL.<sup>7</sup>

What determines the essential aim of a communicative act? Stalnaker has sometimes described assertions as "proposals" to change the context, but without offering a theory of what a proposal is (e.g. Stalnaker 1999, 10–11). In some places, he fills this lacuna by appealing to the intentions with which an act is performed: "In

<sup>5</sup> The most explicit early statements of this view are by Stalnaker (1978) and Gazdar (1981), though it is also tempting to read the view into Carlson (1982), Hamblin (1971), Heim (1982) and Lewis (1979).

<sup>6</sup> I focus on Roberts' influential version of these ideas because she clearly signals her allegiance to a context-directed theory of communicative acts. Many aspects of her theory of questions can also be found elsewhere: Carlson (1982), Ginzburg (1994) and Hamblin (1971).

<sup>7</sup> Other context-directed theories of directives include Condoravdi and Lauer (2012), von Stechow and Iatridou (2017), Roberts (2004, 2015, 2018) and Starr (2019).

my account of assertion, 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). This version of Stalnaker’s theory fits into the intentionalist tradition, along with Grice’s addressee-directed view and other context-directed views that are explicitly framed in terms of intentions to change the context (e.g. Roberts 2012, 2018; Thomason 1990). However, Stalnaker has not consistently endorsed intentionalism. In one place, he denies that an assertion must be intended to make its content common ground, because a speaker may assert *p* despite knowing that *p* won’t be commonly accepted (Stalnaker 1999, 87). Even in these cases, however, he maintains that adding *p* to the context is still the “essential effect” of asserting *p*. His reason is that the speaker’s ability to accomplish their goals with this kind of atypical assertion must be explained “partly in terms of the fact that it would have had certain essential effects had it not been rejected” (Stalnaker 1999, 87). In what follows, I will mostly proceed as if Stalnaker and other context-directed theorists are intentionalists, but this is not essential to my argument. All that matters is that context-directed theorists take updating the context to be the essential aim of a communicative act in the following sense: if an act *were* to achieve this aim (perhaps contrary to expectations), this would make it successful qua the communicative act that it is. To take Stalnaker’s example, if telling O’Leary that he is a fool were to convince him, this would be surprising, but it would also constitute a kind of success. In situations where this is not the case, we should say that my speech act isn’t an assertion at all, but rather an indirect speech act of some kind or a non-communicative attempt at manipulation.

The word ‘essential’ in (CD) is important. Most communicative acts are aimed at many outcomes. In telling you that I am tired, I may aim to inform you that I am tired, to motivate you to leave my apartment, to make you realize that you have overstayed your welcome, to deter you from doing so in the future, and to get some rest. Only the first of these is plausibly the essential aim my assertion—the aim that makes it the kind of communicative act that it is. (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 acts’ essential and secondary 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.<sup>8</sup> Non-intentionalist theories that ground the properties of communicative acts in their essential aims need some other way of drawing a distinction between essential and inessential aims.

<sup>8</sup> Having a communicative intention requires intending to produce an effect in one’s addressee and also intending that they recognize the intention to produce this effect. (Many variations on this idea have been proposed.) Successful communication (i.e. uptake) occurs when the addressee recognizes what effect the speaker intends to have on them. I discuss the different ways that a communicative intention can be fulfilled in Sect. 4.

The other crucial element of context-directed theories, as I am understanding them, is the theory of context that they rely on.<sup>9</sup> The following two principles spell out two key aspects of this theory.

**(CP)** CONTEXT PUBLICITY

Contexts are public: the information in the context is equally available to all interlocutors.

**(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 the interpretation of communicative acts. 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, knowing this, a speaker who wishes to refer to him might do so by uttering ‘Abe’. This example illustrates the idea that the common ground is what hearers rely on as a supplement when linguistic convention does not fully specify what a speaker is saying, and what speakers assume hearers will rely on when designing utterances. Common ground can reliably play this role only insofar as 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.

The idea that context is both the target of communicative acts and also the body of information that we use to interpret them has animated an influential view of anaphora. On this view, an anaphoric expression is one whose use presupposes that the context is in a certain state—a state that determines how the expression will be interpreted. Part of the job of an anaphor’s antecedent is to put the context into a state that satisfies the anaphor’s presupposition. This idea has taken many more precise forms, and has been central to the most common interpretation of dynamic models of conversation. According to Heim’s (1982, 1983) influential theory, for example, a definite NP presupposes that the context entails the existence of a salient individual that meets the NP’s descriptive content, and part of the function of an indefinite NP is to put the context into a state that will satisfy anaphoric definites’ presuppositions. On Heim’s view, the context is more richly structured than Stalnaker’s common ground, since it is structured around discourse referents rather than mere propositions. But Heim makes it clear that she takes her theory to be an extension of Stalnaker’s (Heim 1982, §3.1.4; 1992, 185), and others have followed her in this interpretation (e.g. Roberts 2003). Although it is possible to interpret the formalism of Heim’s and other dynamic theories of anaphora in terms that don’t

<sup>9</sup> Contexts, 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 2019, 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 2018). I gloss over most of these subtleties here.

entail the publicity of contexts or a context-directed theory of communicative acts, the latter views are features of the most natural and widespread interpretations. The popularity of dynamic theories of anaphora therefore partly explains the popularity of context-directed theories of communicative acts.

I am skeptical that we should think of contexts as playing an essential role in linking anaphora to their antecedents, or in the interpretation of communicative acts more generally. That they don't will be a secondary thesis of Sects. 3–5. However, the idea that context often or always plays a role in the interpretation of communicative acts is compatible with my central thesis in this essay, which is that contexts do not play the conceptually distinct role of being the essential targets of communicative acts.

Stalnaker's way of making sense of the idea that contexts are public also entails context reducibility.<sup>10</sup> On his view, for a proposition  $p$  to be in the common ground is for the interlocutors to 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.<sup>11</sup> 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  $p$  isn't in the common ground. It is impossible, by virtue of the nature of common ground, for interlocutors to have unequal access to it.<sup>12</sup>

The reducibility of contexts is an attractive idea for several reasons. It is a given that agents have private propositional attitudes; it would be nice not to treat contexts as further theoretical primitives. More importantly, the psychological reducibility of contexts explains both the publicity of contexts and the roles that they are held to play in explanations of communication. A (truthful) assertion that succeeds in its aims provides the addressee with new information. If the aim of assertion is to change the common ground, then the informativeness of successful assertion is explained only if changing the common ground entails changing the addressee's attitudes. Likewise, some of the attraction of context-directed theories 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

<sup>10</sup> Stalnaker's theory of common ground has evolved over time (Stalnaker 1973, 1978, 1984, 1998, 2002, 2014). I rely on the most recent iteration of his view here (Stalnaker 2014).

<sup>11</sup> Stalnaker thinks of acceptance as a "generic propositional attitude concept" that includes belief and various other more tentative belief-like attitudes, such as supposition, that count as correct only when their contents are true (1984, 79–84). This makes common belief and common knowledge special cases of common acceptance, for Stalnaker.

<sup>12</sup> It is possible to weaken this theory of context in order to forego reducibility while preserving publicity. For example, one could say that each fact about context is reducible to the truth of certain propositions together with the participants' common acceptance of these propositions. This formulation is entailed by Stalnaker's account, on which the propositions in question are themselves about what the participants accept, but it is also compatible with the possibility that the propositions in question aren't all about participants' mental states. All of my objections to context-directed theories later in this essay are unaffected by this weakened view of context. But since I know of nobody who has defended this view, I will treat the stronger version as my target.

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.<sup>13</sup>

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 in a way that is influenced by the state of the TDL. These roles are analogous to the interpretive roles that Stalnaker assigns to common ground. In order to reliably play them, 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 makes them public. Roberts and Portner seem to agree. In answering a question, Roberts (2012, 6, fn. 7) 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”. Likewise, Portner (2004, 242) 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”. The question of how the QUD and TDL are constructed from interlocutors’ private mental states is an interesting and open one: neither Roberts, nor Portner, nor (as far as I know) anyone else has articulated reductions of the kind 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.

### 3 Communication without context change

I turn now to my first objection to context-directed theories. It rests on the observation that it is possible to succeed in communicating with someone without making public either one’s message or the fact that it has been sent, and without having aimed at either of these kinds of publicity. 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 understand, in order for communication

---

<sup>13</sup> The main alternative is to hold that the state of the context is fully determined by grammatical conventions together with facts about what has been uttered so far in the conversation (DeVault and Stone 2006; Lepore and Stone 2015; Stojnić et al. 2017). I don’t have space to adequately discuss this account, and the argument in this paper should be understood as being aimed at context-directed theories paired with a psychologistic theory of context. But one challenge for the view arises from the fact that it does not guarantee the publicity of contexts, since the participants in a conversation may be mistaken about the history of their own conversation. In such cases, it will be what the interlocutors’ accept, rather than the objective state of the context, that determines how the conversation will continue.

to happen. It follows that changing the context can't be the essential aim of communicative acts.

I begin with a scenario modeled on the coordinated-attack problem.<sup>14</sup> Although this scenario is artificial, it has the virtue of clearly illustrating the features of certain real-life conversations that interest me. 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 are certain 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 it is always likely that a messenger will be captured. Suppose that General A sends the message 'we will attack 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 this situation by replying, 'I got your message. We will attack at dawn'. This time, even if the message reaches A, B doesn't know this, and so won't attack, and A will infer B's doubts, and so won't attack either. This process can be iterated indefinitely without solving the problem. Given plausible assumptions about the generals, it can be proved both that they won't attack unless they commonly know that they will attack at the same time, and that they can never achieve common knowledge through their unreliable channel.<sup>15</sup> Nor can the generals 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, they are unable to add information about the time of their attack to the common ground.

Does this show that the two generals can't communicate? On the contrary. 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 and understands a message, successful communication has taken place. Making the time of their attack common ground is a further, futile goal that they have—one that goes beyond the goal of communicating and that they intend to accomplish *by means of* communicating.

To see that this is the right description, suppose that after exchanging several messages with B, A realizes 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 this situation will never be able to coordinate their attack, so they should give up. While he's writing, A adds, he has a question for B:

<sup>14</sup> Some 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) and Yemini and Cohen (1979). The closest precursor of my discussion is Jankovic (2014), who uses the problem to object to theories of communicative acts that are similar to my targets [e.g., those of Schiffer (1972) and Sperber and Wilson (1995)], though for rather different ends.

<sup>15</sup> For proofs, see, e.g., Moses (1986), Yemini and Cohen (1979), and Fagin et al. (1995, ch.6).



- (1) Some of my men keep getting cholera. What should I 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 them to stop eating and drinking near the latrines.

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

Were this scenario to play out, A would have communicated by asking a question with (1) and B would have communicated by performing an assertion (or perhaps a directive) with (2). We must classify these as episodes of successful communication in order to explain their subsequent actions: A prompted B to respond in a highly specific way, and B prompted A to take specific actions to deal with the cholera outbreak. This sort of reliable and precise information sharing is precisely the phenomenon that a theory of communication must explain. But this shows 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 to achieve one's further, extra-communicative aims. From the fact that I don't succeed in getting you to leave my apartment by saying 'I'm tired', it doesn't follow that I haven't communicated 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 extra-communicative aim of gaining common knowledge in order to coordinate their attack.

Furthermore, (1)–(2) show that one can perform a communicative act and even successfully thereby communicate without even *aiming* to make anything public. 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 at this stage. What B cares about in uttering (2) is for A to *believe* the content of his assertion, since having that belief will be sufficient for B to solve his cholera problem. It is unnatural and unmotivated to describe B as intending (or proposing) to change the context.

A final noteworthy observation is that it is possible for the generals to felicitously utter expressions that are anaphoric on earlier utterances whose contents have not been made public and therefore haven't updated the context. In (2), for example, B's use of 'them' is anaphoric on A's use of 'Some of my men' in (1). As I outlined in Sect. 2, this sort of anaphoric link is often explained by saying that indefinites function to update contexts into states that are presupposed by anaphoric definites. B's utterance of 'them' is a counterexample to theories of this kind: it is felicitous despite the fact that its antecedent was uttered in a situation that made it impossible to update the context.

Here is a tempting response to the foregoing line of thought. When an utterance presupposes that the context is in a state that it isn't in, the interlocutors will often *accommodate* this presupposition by silently updating the context to the relevant

state and proceeding accordingly (Lewis 1979). So although general A cannot update the context with the information that some of his men have come down with cholera by uttering (1), perhaps this information is accommodated into the context once general B utters (2) (or once A receives it).

But the presupposition triggered by (2) needn't be accommodated in order for (2) to be felicitous. Suppose that A and B send no more messages after (1)–(2), but see each other in person a short time later, having abandoned their ambitions to attack. It would be natural for their conversation to pick up with B saying the following to A:

(3) Did you receive my answer to your question?

This utterance does not make sense if (2) has updated the context or its presuppositions have been accommodated. Questions like (3) are standardly understood as strategies for repairing the context with the update from an earlier act.<sup>16</sup> Plausibly, this is the point of (3)—i.e. to update the context with the content of (2), since this wasn't possible earlier. In general, the use of a repair strategy like this is a giveaway that the content of an earlier message has not updated the context and its presuppositions have not been accommodated. Of course, repair utterances like (3) often occur immediately after the utterance whose update they are intended to repair, and this can obscure the fact that the original utterance can result in successful communication before the repair takes place. What is different about this case is that general B has already successfully communicated with A before A has a chance to repair the context—a conclusion that must be drawn in order to explain A's actions after (2) but before (3).

It's not hard to see why the generals wouldn't accommodate the presupposition of (2), if accommodating a presupposition entails adding it to the context and therefore making it public. The generals cannot make information public via accommodation any more than via assertion. If they could, then their coordination problem could be solved by uttering something that presupposes that the attack will be at dawn—e.g., 'The attack at dawn will no doubt succeed'—and accommodating this presupposition. Even after B utters (2), B won't accept that A has received it, and so won't accept that A has accommodated its presupposition. And when A receives (2), A will infer this about B, and so won't accept that B accepts that A accepts its presupposition.

The coordinated-attack scenario is admittedly an artificial one for the purposes of studying human communication. It may be tempting to reply that my argument affects only certain bizarre or hypothetical edge cases, and that the proper response is to ignore them or idealize them away until we have a solid theory of the core cases of human communication. But this response is mistaken. Many actual cases of successful communication are relevantly analogous to the coordinated-attack scenario, in that they involve a speaker who can't make the content of their communicative act public and knows this, but who succeeds in their communicative

<sup>16</sup> Repair strategies of this kind are sometimes called 'grounding' or 'metacommunication' (Clark 1996; Clark and Brennan 1991; Clark and Schaefer 1989; Ginzburg 2012).

aims anyway. I will call situations of this kind *publicity averse*. Publicity-averse situations are both quite ordinary and also highly relevant to the topic of this essay, since they cast the distinction between the essential and inessential aims of communicative acts in sharpest relief.

One particularly clear sort of publicity-averse situation arises whenever someone performs a communicative act that they intend to be understood only after they're dead. Suicide notes and wills have this feature. Likewise, suppose that I write a letter to my daughter that, I stipulate, should not be opened until after my death. I fully expect—indeed, I *intend*—that my message won't reach my daughter until I am no longer around to partake in propositional attitudes. If everything goes according to plan, the content of my letter will never be common ground for me and my daughter. If it did become common ground, this would constitute failure rather than success. The communicative acts I perform with my letter therefore can't be thought of as proposals to update the common ground. Still, there is nothing far-fetched about communicating via such a letter.

A distressingly common kind of publicity-averse situation occurs whenever one sends an email to a recipient with an overwhelmed inbox.<sup>17</sup> I have friends with unconscionably high unread-message counts, whose likelihood of actually opening and reading any given email they receive is surely as low as the chance of a messenger making it through an enemy-infested valley. When I assert something in an email to one of these acquaintances, I continue to doubt that they have read it until I encounter evidence to the contrary, and so my aim in emailing them normally won't be to add to the common ground. But it may still be perfectly reasonable for me to email them with the communicative aim of changing their beliefs, and I sometimes succeed in this aim without knowing it.

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 sometimes 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 (4).

(4) 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:

(5) ME: Did you read my email?

(6) FRIEND: Yes, I read it.

(7) ME: The conference is cancelled, so I won't have to leave early after all.

Although I might have led with (7), it is also natural for me to first begin with (5), particularly if I prefer to minimize the chance of misunderstanding. But (5) is the sort of repair strategy I mentioned earlier, and makes no sense if the content of my email is

<sup>17</sup> The point that email exchanges can be relevantly similar to the coordinated-attack scenario has been discussed in the context of the 'electronic mail game' (Rubinstein 1989; Lederman 2018). Jankovic (2014) discusses this point in a way that parallels my argument here.

already common ground. However, the fact that merely confirming that my friend has received the email can suffice for this kind of repair demonstrates that he has already understood my original message. 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 for that content to make its way into the common ground. We can also suppose that my friend had taken various actions as a result of understanding (4), but before I uttered (5). Suppose that he had moved topics relevant to me to the beginning of his lesson plan, for example. These actions would have to be explained by saying that we'd already communicated before my utterance of (5) could repair the context.

This is another case of communicating  $p$  without making  $p$  (or the fact that one has asserted  $p$ ) common ground. Moreover, I see no good reason to think of my assertion as a proposal to add  $p$  to the common ground. I fully expected not to do so, given my knowledge of my friend's inbox. And again, it is plausible that all I care about, in this case, is that my friend *believe* that I have a good reason for leaving his class. My purpose for caring about this would be no better served by making my reason for leaving common ground.

Publicity-averse scenarios like this are boringly common. 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 places where English is not widely spoken, 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've 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 teaching a class or giving a talk I sometimes use technical vocabulary that I believe some of the audience members won't understand. Sometimes it turns out that they did understand me, in spite of my doubts.

We often immediately resort to context-repair strategies in situations like these. But this follow-up normally happens *after* the hearer has understood the speaker's message. All of these situations are temporarily publicity averse, if only for a short time. This becomes clear when we consider situations, such as the coordinated-attack scenario, in which there is a delay between the initial communicative act and the repair that is long enough for the addressee to begin acting on the information communicated. When an addressee performs actions on the basis of specific information gained as a result of having understood a communicative act, I take this to be a strong reason to conclude that communication has occurred, and I take this to be the sort of event that a theory of communication needs to explain. Each of the publicity-averse scenarios I have discussed have this feature: General B stops his cholera outbreak as a result of having understood General A's message, executors of wills take complex action on the basis of what they learn from the will, my friend changes his lesson plan on the basis of having understood my email, and so on.

A final important point about real-world, publicity-averse situations is that, like the coordinated-attack scenario, they sometimes allow for the felicitous utterance of anaphora whose antecedents haven't updated the context. Suppose that a student who has visited my office hours rarely understands anything I say, and today is also wearing headphones and seeming inattentive. Thinking that it is possible but unlikely that I will succeed, I nonetheless proceed to explain some concepts that the student has sought help about, uttering (8).

- (8) An important axiom is that two sets are identical if and only if they have the same members. **This axiom** is very important. **It** entails that no two sets can have the same extension. ...You must understand **the axiom** in order to pass this class.

Each of the boldfaced definite NPs in (8) is anaphoric on the indefinite, 'an important axiom'. According to the popular theory of anaphora that I sketched in Sect. 2, these anaphoric links are made possible because the context was updated by 'an important axiom' so that it entailed the existence of a salient axiom, and definites I utter later presuppose that the context is in this state, which in turn determines how they will be interpreted. But if contexts are bodies of information that are necessarily public, then this can't be the right way to think about (8). My best guess when uttering all of (8) is that my student has no idea what I'm talking about, if they can even hear me at all. And so, when I utter the later definites, the context, taken as a public body of information, does not meet their presuppositions and cannot do the job of determining how they should be interpreted. The possibility that these presuppositions haven't been accommodated by the time I've finished uttering (8), and the fact that I have nonetheless succeeded in communicating its content to my student, can be seen from the fact that the following would be a natural way for the conversation to proceed:

- (9) ME: Did you hear a word of what I said just now?  
 STUDENT: Sure I did. By the way, am I right that it follows from the axiom that there can be only one empty set?

From the considerations adduced in this section, I conclude that communicating requires changing the addressee's private mental state, but not changing contexts, understood as public bodies of information. Although it may be that communicative acts are sometimes performed with the aim of changing the context, this aim an optional, extra-communicative goal that we may fail at without failing to communicate, and that therefore isn't part of the essential aim of performing a communicative act.

#### 4 Ways that a communicative act can succeed

My plan in this section is to distinguish seven distinct ways that a communicative act can succeed according to a simple version of addressee-directed intentionalism. Whereas addressee-directed intentionalists can correctly predict these to be real distinctions, context-directed theories elide several of them.

I will focus on a genus of communicative acts that is unusually broad from the point of view most addressee-directed theorists. 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 but also various more provisional belief-like attitudes that don't entail belief. Assertion, as Grice thinks of it, has the more precise aim of producing belief. Stalnaker sometimes acknowledges a parallel consideration by saying that assertion, as he defines it, is a broader notion than the one that is often discussed by philosophers (Stalnaker 2018). 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 view, 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 entails the communicatively intended effect of assertion\* as I have defined it. I will ignore non-assertoric acts 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 an addressee-directed intentionalist, are seven ways in which such an assertion\* could succeed, along with corresponding success conditions.<sup>18</sup>

- (S1) PERFORMANCE  
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<sup>19</sup>  
A accepts that *S* uttered *x* intending for *A* to accept *p*.
- (S3) ACCEPTANCE  
A accepts *p*.
- (S4) INFLUENCE  
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*.

<sup>18</sup> I do not mean to suggest that this list is exhaustive.

<sup>19</sup> Some critics of Grice have argued that his theory must be altered to reflect the fact that uptake, rather than acceptance, is the primary aim of communication (Strawson 1964; Searle 1969; Bach and Harnish 1979). Grice's (1989, 352) response is to point out that it had been part of his ambition to analyze uptake as the addressee's successful recognition of the speaker's intention to produce a response—i.e., as what I am here calling uptake (S2). On Grice's view, performing a communicative act (i.e., meaning something by an utterance) always requires intending to produce both a response in an addressee (S3) and uptake (S2), and the latter form of success is the criterion of successful communication.

## (S7) ANAPHORA LICENSING

Utterances anaphoric on  $x$  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 sometimes come up in discussions of rejected speech acts.<sup>20</sup>

Addressee-directed intentionalism predicts that each of these kinds of success is distinct. I contend that this is a good prediction: we really do need to draw all of these distinctions in order to make sense of how and why humans communicate. And, as I began to argue in Sect. 3, a theory of communication should center around S1–S3 rather than S5–S6.

First, with the possible exception of S6 being sufficient for S7, none of these forms of success entails the subsequent forms. One can perform a communicative act without being understood in any way: S1 does not invariably lead to any of S2–S7. After all: the thing that an addressee fails to understand in such a case is a communicative act. Likewise, S2 does not invariably lead to S3, nor S3 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. Analogously, S5 does not invariably lead to S6: even when it is commonly accepted that a certain communicative act has been performed, they act may still be rejected.

The interesting questions, for my purposes, are about the relationships between the context-free forms of success, S2–S4, on one hand, and the context-involving forms, S5–S6, on the other. My arguments in Sect. 3 were designed to show that S2–S4 can take place without either S5 or S6, that performing a communicative act requires aiming at S3 but not at S6, that successful communication may involve S2 but not S5–S6, and that S7 may occur without S5–S6. In a publicity-averse situation, a speaker cannot hope to achieve common uptake or common acceptance. They *can*, however, achieve uptake, acceptance, and influence, and they can sometimes even thereby license downstream anaphora. Moreover, in each of the publicity-averse scenarios I've discussed, the speaker cares more about achieving S2–S4 than about S5–S6. In some of the scenarios, it is unnatural to say that the speaker's aims include S5–S6 at all.<sup>21</sup>

<sup>20</sup> Clark and Schaefer (1989) and Ginzburg (2012) consider cases in which S5 and S6 fail, but treat this phenomenon only as a kind of failure to communicate. They fail to distinguish S1–S4 as distinctive kinds of success that may be valuable on their own, and so do not consider cases in which speakers either don't aim for or don't care about S5–S6.

<sup>21</sup> Stalnaker (1999, 87) thinks that a speaker can assert  $p$  knowing that they can't add  $p$  to the common ground. He gives the example of telling O'Leary that he is a fool, presumably with little hope of convincing him. In such cases, Stalnaker thinks, the speaker's "primary intention" is to achieve some "secondary effect" rather than the assertion's "essential effect". Still, Stalnaker argues that we should think of adding  $p$  to the common ground as the essential effect of assertion, even in these cases, because "one generally explains why the action has the secondary effects it has partly in terms of the fact that it would have had certain essential effects had it not been rejected".

The cases I have discussed are different from those discussed by Stalnaker, since they don't involve communicative acts getting rejected. Still, I anticipate his line of thought inspiring the following objection: S6 (common acceptance) is the essential effect of assertion\*; however, in publicity-averse

On the other hand, S6 entails S3 and S5 entails S2: it is impossible for S and A to commonly accept something unless A accepts it. This shows that the kinds of success on which context-directed theorists focus are special cases of kinds of success on which addressee-directed theorists focus. Addressee-directed theories are thus more general and account for a wider range of cases of communication than context-directed theories.

Finally, consider anaphora licensing (S7). In Sect. 3, I argued, contrary to a widespread interpretation of dynamic theories of anaphora, that S7 does not require S5 or S6: an anaphoric expression can be felicitous even if its antecedent has not been used to update the context. Strictly speaking, this claim is independent of my central thesis in this essay: even if context-directed theories are wrong about the nature of communicative acts, it could still be the case that updating the context is necessary in order to license anaphora (and vice versa). The fact that it is not necessary, however, further undermines the idea that context updates should be granted a central role in a theory of communicative acts, since it undermines the idea that context updating plays one of the most important roles that it has been thought to play in semantics and pragmatics.

In particular, it might have been tempting to reply to my arguments in this essay as follows. “Let us stipulate that there is a sense of ‘communication’ on which the inhabitants of publicity-averse scenarios can communicate. Label it ‘communication<sub>1</sub>’. Still, there may be another sort of activity, communication<sub>2</sub>, that is of theoretical interest to semantics or pragmatics, and that cannot take place in publicity-averse scenarios because it entails S5–S6.

It is true that I have not ruled out this possibility, though my argument that anaphora licensing does not require context update rules out one role that we might have wanted communication<sub>2</sub> to play. This essay may be read as a challenge to identify a role that communication<sub>2</sub> *does* invariably play. (Of course, even if we found such a role, this would not undermine the importance of communication<sub>1</sub>).

It may be asked: what *is* necessary for anaphora licensing if not common acceptance (S6) or even common uptake (S5)? I won’t attempt to articulate a full answer to this question, but here are some preliminary thoughts.

In Sect. 3, I gave two examples of felicitous anaphoric links in publicity-averse contexts. In both examples, achieving acceptance (S3) requires the hearer to interpret the speaker’s utterance so as to incrementally add to their information about an entity that they already represent.<sup>22</sup> One part of this task is to infer *which*

---

Footnote 21 continued

scenarios, we sometimes assert something knowing that we won’t achieve this effect but hoping that we will achieve a secondary effect instead—namely, S3 (acceptance). My reply is that I take myself to have given good reasons for thinking that it is S6 that should be considered a secondary effect—one that is therefore more demanding than S3, one that speakers care about producing with some communicative acts but not others, one at which they needn’t aim in order to make assertions, to successfully communicate, or to license downstream anaphora, and one whose achievement would sometimes constitute failure rather than success.

<sup>22</sup> My wording is intended to leave open the possibility that this incrementation of information about previously represented entities can be modeled as the filing of new information under an old discourse referent, in keeping with dynamic models of discourse. But whereas discourse referents are often thought



previously represented entities the new information should be linked to. General A must infer that it is *his troops* whom general B is saying should stop dining near the latrines, for example. And my student must infer that it is *the axiom I described at the start of my office hours* that I am saying he must understand in order to pass. An important assumption of the kind of theories I have been criticizing is that hearers infer anaphoric connections of this kind by drawing only on information that is entailed by the public context. But of course, any given interlocutor will also possess plenty of information that is not contextually entailed. What I take myself to have shown is that hearers can at least sometimes discern anaphoric links and thereby increment their information as intended by drawing on their private, non-contextual information.

Here, then, is a tentative suggestion about what is required in order for an utterance of an anaphoric expression,  $x$ , to be felicitous: (i) the addressee either accepts the presuppositions that  $x$  triggers or is willing to change what they accept accordingly and without comment, and (ii) the addressee is capable of inferring which previously represented entity is at issue in order to increment their information as intended by the speaker, using any (possibly-non-contextual) information at their disposal. These two conditions are entailed by standard public-context-centric interpretations of dynamic theories, but are strictly weaker than those theories, since they don't give information entailed by the public context an essential role to play.

This view explains why some uses of anaphoric expressions are infelicitous. Suppose, for example, that someone begins a conversation by uttering (10):

(10) She is very smart.

It is normally infelicitous to utter a pronoun discourse-initially and in the absence of a salient referent. A standard explanation for this is that such uses of pronouns trigger presuppositions that cannot be accommodated because interlocutors are unable to infer *how* to update the context in order to accommodate them. My explanation is similar: in trying to interpret (10), the hearer may have no way of inferring which previously represented entity it is whom the speaker is saying is smart, and so will be unable to increment their information as the speaker intends.

## 5 Context-directed theories and idealized models

Before concluding, I want to address a potential reply to my arguments. Context-directed theories, this reply goes, are normally components of idealized models of conversation. As purveyors of idealized models, context-directed theorists do not (or at least should not) claim that they can explain every nuance or instance of

---

Footnote 22 continued

of as features of public contexts, I have in mind features of hearers' private mental states. The idea of modeling mental content in ways that are analogous to the contextual structures posited by dynamic theorists has been explored by those who posit 'mental dossiers' (Forbes 1990), 'mental files' (Recanati 2012; Huemer et al. 2018), and 'discourse content' (Cumming 2014).

communication. Rather, their models 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 their explanatory goals. 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 it is correct that context-directed theories are most charitably understood as components of idealized models, but we must be cautious about what follows from this. It is legitimate to draw some conclusions but not others from idealized models. Consider a pair of standard examples that have been discussed by philosophers of science. Galileo's model of projectile motion, which ignored the effects of air resistance, was useful for understanding the effects of gravity on projectiles, but would not have helped us to understand the effects of air resistance (Weisberg 2007; McMullin 1985). Likewise, explanations of Boyle's law that idealize away particle collisions in low-pressure gasses proved useful for understanding the macroscopic behavior of gasses, but would not have been useful for understanding the details of their microscopic behavior (Weisberg 2007; Strevens 2011). The lesson is that we must be careful to distinguish distortions introduced by an idealized model from features of the target phenomenon being modeled. My claim is that theories that give public contexts an essential role to play in understanding communicative acts and anaphora have failed to do this.

In the real world, a speaker's grasp of their addressee's mental state is always partial and imperfect. This includes the speaker's grasp of the information that their hearer will use to interpret their utterances, as well as their grasp of the mental states that they seek to change. To speak is therefore always to incur risk: taking too much for granted may leave the addressee unable to understand, but taking too little for granted will result in wasting their time. The question of how speakers manage to steer between these hazards is a very difficult one. Answering it with any detail will require a much fuller understanding of the human capacity for mindreading than cognitive science can currently offer us. By constructing idealized models of conversation in which speaker and hearer make use of the same body of information, and in which this is also the body of information to which they seek to add with their communicative acts, we idealize away from an aspect of conversation that we currently lack the means to understand. The hope, I take it, is that by modeling conversations as abstract processes in which interlocutors draw on and add to a single, well-defined body of information, thereby idealizing away from the role of mindreading, we can tease out some grammatical, or quasi-grammatical principles at the core of a messy and theoretically intractable social practice. Nothing I have said in this essay shows this theoretical aim to be misguided.

However, when working with dynamic models of conversation, we must distinguish aspects of the model that reflect generalizations about actual language use from distortions introduced by the model itself. If my arguments in this essay have been sound, then the idea that public, intersubjective contexts invariably serve as the targets of communicative acts or as the link between anaphora and their antecedents is, at best, an artifact of idealized models. To mistake these stipulative features of models for empirical generalizations about real conversations would be

like concluding, from Galileo's model, that a lack of air resistance is an essential feature of projectile motion.

## 6 Conclusion

Contrary to current trends, communicative acts should not be understood as attempts to pool information with one's peers. On the contrary: communicative acts needn't aim at affecting the conversation's public context in any way, and communication can be successful without changing the context. This is even true in at least some cases when later stages of a conversation build on former stages by means of anaphoric links.

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 if communication sometimes does result in context change, this should be understood as an extra-communicative accomplishment, much as offending or prompting someone to action are extra-communicative accomplishments.

**Acknowledgements** I have received helpful feedback from Kent Bach, Anne Bezuidenhout, Justin Bledin, Elisabeth Camp, Kai von Fintel, Daniel Fogal, Samia Hesni, Ernie Lepore, Matt Mandelkern, Eliot Michaelson, Sarah Murray, Stephen Neale, Gary Ostertag, Craige Roberts, Will Starr, Andreas Stokke, Elmar Unnsteinsson, Seth Yalcin, audiences at Rutgers and the University of Oslo, and two anonymous reviewers for *Philosophical Studies*.

## References

- Akkoyunlu, E. A., Ekanadham, K., & Huber, R. V. (1975). Some constraints and tradeoffs in the design of network communications. *ACM SIGOPS Operating Systems Review*, 9(5), 67–74.
- Austin, J. L. (1962). *How to do things with words*. Oxford: The Clarendon Press.
- Bach, K., & Harnish, R. M. (1979). *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.
- Camp, E. (2018). Insinuation, common ground, and the conversational record. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts*. Oxford: Oxford University Press.
- Carlson, L. (1982). *Dialogue games: An approach to discourse analysis*, Number 17 in Synthese Language Library Dordrecht: D. Reidel.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. Resnick, B. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). Washington, DC: American Psychological Association.
- Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, 13(2), 259–294.
- Cohen, D., & Yemini, Y. (1979). Protocols for dating coordination. In *Proceedings of the fourth Berkeley conference on distributed data management and computer networks* (pp. 179–188). Berkeley, CA: Lawrence Berkeley Laboratory.
- Condoravdi, C., & Lauer, S. (2012). Imperatives: Meaning and illocutionary force. In C. Pinõn (Ed.), *Empirical issues in syntax and semantics 9: Papers from the Colloque de Syntaxe et Sémantique à Paris 2011* (pp. 37–58).
- Cumming, S. (2014). *Discourse content* (pp. 214–230). Oxford: Oxford University Press.
- DeVault, D., & Stone, M. (2006). Scorekeeping in an uncertain language game. In D. Schlangen, & R. Fernández, (Eds.), *Brandial 2006: Proceedings of the 10th workshop on the semantics and pragmatics of dialogue* (Vol. 10, pp. 139–146). SemDial, University of Potsdam.

- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about knowledge*. Cambridge, MA: MIT Press.
- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1999). Common knowledge revisited. *Annals of Pure and Applied Logic*, 96(1), 89–105.
- Forbes, G. (1990). The indispensability of sinn. *The Philosophical Review*, 99(4), 535–563.
- Gazdar, G. (1981). Speech act assignment. In A. K. Joshi, B. L. Weber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 64–83). Cambridge: Cambridge University Press.
- Ginzburg, J. (1994). An update semantics for dialogue. In *First international workshop on computational semantics (IWCS 1)* (pp. 111–120).
- Ginzburg, J. (2012). *The interactive stance: Meaning for conversation*. Oxford: Oxford University Press.
- Gray, J. (1978). Notes on data base operating systems. In R. Bayer, R. M. Graham, & G. Seemüller (Eds.), *Lecture notes in computer science 60: Operating systems* (pp. 394–481). Berlin: Springer.
- 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, P. (1989). Retrospective epilogue. *Studies in the way of words* (pp. 339–385). Cambridge, MA: Harvard University Press.
- Halpern, J. Y., & 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.
- Harris, D. W., Fogal, D., & Moss, M. (2018). Speech acts: The contemporary theoretical landscape. In D. Fogal, D. W. Harris, & M. Moss (Eds.), *New work on speech acts*. Oxford: Oxford University Press.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. Ph.D. Dissertation, University of Massachusetts at Amherst.
- Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, use and interpretation of language* (pp. 164–189). Berlin: de Gruyter.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3), 183–221.
- Huemer, M., Perner, J., & Leahey, B. (2018). Mental files theory of mind: When do children consider agents acquainted with different object identities? *Cognition*, 171, 122–129.
- Jankovic, M. (2014). Communication and shared information. *Philosophical Studies*, 169(3), 489–508.
- Lederman, H. (2018). Two paradoxes of common knowledge: Coordinated attack and electronic mail. *Noûs*, 54(4), 921–945.
- Lepore, E., & Stone, M. (2015). *Imagination and convention*. Oxford: Oxford University Press.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3), 339–359.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science*, 16(3), 247–273.
- Moses, Y. (1986). *Knowledge in a distributed environment*. Ph.D. thesis, Stanford, CA.
- Portner, P. (2004). The semantics of imperatives within a theory of clause-types. In K. Watanabe, & R. Young (Eds.), *Proceedings of SALT 14*. CLC Publications.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15, 351–383.
- Recanati, F. (2012). *Mental files*. Oxford: Oxford University Press.
- Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26, 287–350.
- Roberts, C. (2004). Context in dynamic interpretation. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 197–220). Oxford: 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 T. Brochhagen, F. Roelofsen, & N. Theiler (Eds.), *Proceedings of the 20th Amsterdam Colloquium* (pp. 353–362).
- Roberts, C. (2018). Speech acts in discourse context. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts*. Oxford: Oxford University Press.
- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *The American Economic Review*, 79(3), 385–391.
- Schiffer, S. (1972). *Meaning*. Oxford: Oxford University Press.
- Searle, J. (1969). *Speech acts*. London: Cambridge University Press.

- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition*. Oxford: Blackwell.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2, 447–457.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics 9* (pp. 315–332). New York: Academic Press.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- 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: Oxford University Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Stalnaker, R. (2014). *Context. Context and Content*. Oxford: Oxford University Press.
- Stalnaker, R. (2018). Dynamic pragmatics. In D. Fogal, D. W. Harris, & M. Moss (Eds.), *New work on speech acts*. Oxford: Oxford University Press.
- Starr, W. (2019). A preference semantics for imperatives. Unpublished Manuscript.
- Starr, W. B. (2010). *Conditionals, meaning and mood*. Ph.D. thesis, Rutgers University.
- Stojnić, U., Stone, M., & Lepore, E. (2017). Discourse and logical form: Pronouns, attention, and coherence. *Linguistics and Philosophy*, 40(5), 519–547.
- 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*. Cambridge, MA: Harvard University Press.
- Thomason, R. H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication, Chapter 16* (pp. 325–364). Cambridge, MA: MIT Press.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25(3), 221–261.
- von Fintel, K., & Iatridou, S. (2017). A modest proposal for the meaning of imperatives. In A. Arregui, M. Rivero, & A. P. Salanova (Eds.), *Modality across semantic categories, Chapter 13* (pp. 288–319). Oxford: Oxford University Press.
- Weisberg, M. (2007). Three kinds of idealization. *Journal of Philosophy*, 104(12), 639–659.
- Yemini, Y., & Cohen, D. (1979). Some issues in distributed process communication. In *Proceedings of the 1st international conference on distributed computing systems* (pp. 199–203).

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.