

Manipulative Underspecification

Justin D'Ambrosio

Arché Philosophical Research Centre, University of St. Andrews
Forthcoming in *Philosophical Review*

Abstract

In conversation, speakers often felicitously underspecify the content of their speech acts, leaving audiences uncertain about what they mean. This paper discusses how such underspecification and the resulting uncertainty can be used deliberately, and manipulatively, to achieve a range of noncommunicative conversational goals—including minimizing conversational conflict, manufacturing acceptance or perceived agreement, and gaining or bolstering status. I argue that speakers who manipulatively underspecify their speech acts in this way are engaged in a mock speech act that I call *pied piping*. In pied piping, a speaker leaves open a range of interpretations for a speech act while preserving both plausible deniability and plausible assertability; depending on how the audience responds, the speaker can retroactively commit to any of the interpretations left open, and so try to retroactively update the common ground. I go on to develop a model of how pied-piping functions that incorporates game-theoretic elements into the more traditional common-ground framework in order to capture the uncertainty of updating.

Keywords: underspecification, manipulation, strategic speech, common ground, context

Word count: 14,415

1. Introduction

Underspecification is a pervasive feature of natural language communication. Frege famously held that underspecification—and in particular

This paper has benefitted from the help of many people. I particularly grateful to Alan Hájek and Daniel Stoljar for their comments on numerous previous versions, as well as for many helpful discussions of this and related material. Thanks also to Brian Hedden, Wolfgang Schwarz, Martín Abreu Zavaleta, Andrew Peet, Kyle Blumberg, Christian

vagueness¹—was a flaw of natural language, one which, in aspiring to use to a logically perfect language, we should aim to eliminate. But since Frege, much work in pragmatics and philosophy of language has aimed to show how underspecification, rather than being a flaw, is often perfectly rational given the goals of, and the constraints faced by, agents in real communicative situations. For instance, several authors have argued that underspecification is one way for a speaker to navigate the tension between conflicting Gricean constraints, such as the tension between informativity and truth.² Others have pointed out that once we think of conversation as an activity whose goal is efficient communication and recognize that speakers need only specify enough to achieve this goal, underspecification is often the rational choice for a fully cooperative speaker.³

However, speakers do not always have the goal of communicating efficiently. Consider the following situation.

Electoral Strategy I am a politician, you are a member of my constituency, and we meet at a town hall. Your goal is to come to know my policies and plans, and whether you vote for me will be determined by whether you think our positions are aligned. My only goal is to get you to vote for me. You ask me what steps I plan to take concerning gun control, but I have no knowledge

Barry, Derek Ball, Kristina Liefke, Dan Lassiter, Wataru Uegaki, Brian Rabern, Ben Henke, and Zoltán Gendler Szabó. Two referees for the *Philosophical Review* also provided valuable and constructive feedback. Earlier versions of this essay were presented at the Australian National University, the University of St. Andrews, Cambridge University, Ruhr-University Bochum, the University of Stirling, the Australasian Association of Philosophy, and the Edinburgh Meaning Sciences Group.

¹The exact relationship between vagueness and underspecification is a matter of dispute; the phenomena are closely related and not always clearly distinguished—see Williamson (1994, Ch.s 2-3) for discussion. MacFarlane (2020b,c) has recently argued that vagueness is a special case of underspecification, which he in turn takes to be a form of indecision. Supervaluationists such as Lewis (1970, 1986), Fine (1975), and Kamp (1975) likewise tend to treat vagueness as a form of underspecification. But other approaches to vagueness—such as the epistemic approach (Sorensen, 1988; Williamson, 1994; Horwich, 1998) and the metaphysical approach (Tye, 1990; van Inwagen, 1988; Morreau, 2002)—treat them as completely different phenomena. Here I will focus on underspecification, which I take to concern the relation between a speech act and its content, and remain neutral on whether or not it subsumes vagueness.

²For discussion of this view see Horn (1984), Blutner (2000), van Deemter (2010), and Egré et al. (2023).

³See, for instance, Jaegher (2003), O’Connor (2014, 2015), MacFarlane (2016, 2020a,b,c), and King (2018, 2021).

of what your background views on the subject are, and you have no idea whether I intend to be forthright about my plans.

Electoral Strategy is what we can call a *strategic context*: it is a context in which my primary conversational goal is not to communicate or share knowledge with you. Instead, I am aiming to realize a goal that is plausibly in tension with straightforwardly answering your question: I want you to believe that our views are aligned, and so vote for me. But **Electoral Strategy** is also what I call an *impoverished context*. An impoverished context is one in which facts that are crucial for the conversation—facts on which subsequent conversational actions depend—are unknown to one or another interlocutor. In this case, what I should say to achieve my goal depends on facts about you that I don't know: what your background views are. Likewise, how you should interpret me, and whether you should accept my speech acts or dismiss me as a vote-pandering politician, depends on a fact about me that you don't know: whether I am trustworthy. In impoverished contexts, conversational actions depend on reasoning under uncertainty.

Contexts that are strategic and impoverished, in which speaker and audience must reason under uncertainty about each other's conversational goals, are extremely common—perhaps even the norm—in public and political discourse. In such contexts, it is often the case that the best way for a strategic speaker to realize their goals is to deliberately underspecify the content of their speech act. Suppose that in response to your question, I utter (1):

- (1) I will do everything in my power to keep guns out of the hands of dangerous people.

In uttering these words, I deliberately leave open a range of fully specific interpretations for my utterance, and so leave you uncertain about what I mean. Even given the information in the context, the quantificational expression “everything in my power” can be interpreted in a range of ways—I might mean “everything in my power consistent with the wording of the Constitution,” “everything in my power including changing that wording,” or anything in between. Much the same is true for the multidimensional gradable adjective “dangerous,” whose precise extension will not be fixed by context. Given the paucity of contextual information, just as you are uncertain whether I'm trustworthy, you will also be uncertain about the content of my assertion.⁴

⁴The uncertainty of communication is something that has recently received a lot of

Underspecifying in this way affords me several strategic advantages. First, underspecifying makes it difficult for you to challenge or object to my proposal. For each specific content that I might have intended, you will be highly uncertain whether I meant that thing, so you will be uncertain about what to challenge or deny. Moreover, the only content to which you can be certain I am committed is something to which it is nearly impossible to object; who could deny that we ought to keep guns out of the hands of dangerous people? But if my speech act goes unchallenged, underspecification affords me flexibility in what I commit to in downstream conversation. In particular, your uncertainty allows me to plausibly, and retroactively, commit to whatever more specific content is most advantageous to me. Depending on how you respond, I can plausibly claim to have asserted any of the more specific interpretations left open by the context and so attempt to retroactively update the common ground.

Understanding and modelling this form of speech and the contexts in which it occurs requires a number of departures from idealized models of linguistic communication. First, in contexts that are strategic and impoverished, conversational moves of both speaker and audience involve reasoning under uncertainty about how cooperative the speaker is and how trusting the audience is. This requires us to depart from the twin assumptions made by standard frameworks—including the common ground framework—that conversation can be treated as a *cooperative game of complete information*.⁵ An adequate model must allow for the possibility that conversation is a game in which what an audience takes a speaker to have asserted depends on their reasoning under uncertainty about the speaker’s goals. Second, and most crucially, modelling communication in the presence of underspecification requires allowing that audiences can be uncertain about what a speaker asserts, update with the content they find most probable given their pre-suppositions, but revise this update in the course of conversation.

Here, I develop an account of the pragmatics of manipulative underspecification, first by giving an account of the kind of act a speaker undertakes in manipulatively underspecifying, and then by developing a formal model of how this act affects conversation. I first argue that speakers who

attention, perhaps most notably from Peet (2016, 2017), but see also Abreu Zavaleta (2021a,b), Drożdżowicz (2022), and Pollock (2023)

⁵See Stalnaker (1974/2002, 2002, 2014) and Roberts (1996) for discussion of the common-ground framework, which will be discussed extensively below. See Szabó (2020) for discussion of how the common-ground framework treats conversation as a cooperative game in which interlocutors have complete information about each other’s goals.

aim to manipulate by underspecifying engage in a distinctive type of mock speech act that I call *piéd piping*, and I show that *piéd piping* has the effect of temporarily *fragmenting* interlocutors' conceptions of the common ground. I then develop a formal model of how such fragmentation occurs and how it is resolved in downstream conversation. The model integrates a decision-theoretic approach to conversational updating with the more traditional, non-probabilistic approach. Crucially, the new model allows us to capture an audience's uncertainty about both the content of a speech act and whether the speaker is being cooperative in undertaking it. It also reveals the relation between these two forms of uncertainty.

The rest of the paper proceeds as follows. I begin by discussing underspecification in non-strategic conversation and the various theories of how communication functions in its presence (§§2-3). I then discuss examples of how underspecification can be deployed strategically, using these to illustrate the advantages it provides, and I explain why such underspecification falls into the broader category of manipulative speech (§§4-5). I then discuss *piéd piping*, show how it serves to fragment the context, and develop a formal model that captures the dynamics of fragmentation and retroactive update (§§6-7). I conclude by discussing some of the consequences that this model has for our understanding of non-strategic assertion and underspecification (§8).

2. Felicitous Underspecification

Our utterances, on their own, often fail to provide an audience with enough information to determine what we said in speaking as we did. Perhaps the simplest way that this can happen is with ambiguity, either syntactic or semantic. If I utter "The rabbi married my sister," independently of further information, you won't know whether I meant that she was married *by* him or *to* him. What allows you to bridge the gap between my utterance and my intentions is the context. Typically, contexts are treated as bodies of information common, and presumed to be common, among conversational participants, on which speakers rely in planning their utterances and on which audiences rely in inferring a speaker's intentions. Standardly, context will provide you with enough information to resolve any lexical and syntactic ambiguities in my utterance, which in turn allows you to recognize what I intended to say in speaking as I did.

But lexical and syntactic ambiguity are far from the only ways that the linguistic meaning of the words we utter can underdetermine what we say with them. Natural language contains a host of expressions whose semantic

values depend on the context in which they are uttered—expressions that Jeffrey King (2018, 2021) calls *supplementives*. If I utter “We had roast duck last night,” in order to figure out what I’ve said, you need to know (a) who is speaking, (b) who the other salient person or people are that compose the “we,” (c) what day it is, and (d) the fact that in the context, “had” means “ate for dinner.” Only then will you be in a position to know what I said. In this case, context supplements the linguistic meaning of my utterance by fixing the semantic values of its constituent context-sensitive expressions. But in other cases, context may supplement the linguistic meaning of my utterance in ways that are not traceable to any such context-sensitive expressions, as in cases of conversational implicature (Bach, 1994), explicature (Sperber and Wilson, 1986; Carston, 1988), or “free enrichment” (Recanati, 2002).⁶

More recently, however, philosophers of language have begun to recognize that even given contextual information, audiences may not be able to determine a unique semantic value for a supplementive, but that such utterances can nonetheless be pragmatically faultless. To see this, suppose we are discussing Ben, and I utter (2):

(2) Ben is athletic.

In order to determine the semantic value of a gradable adjective—especially one such as “athletic” that is multidimensional—audiences need to rely heavily on contextual information.⁷ Context is needed to determine, at the very least, (1) the relevant dimensions of athletic comparison, (2) the correct way or ways of aggregating or weighting those dimensions, (3) the relevant comparison class, and (4) the degree-threshold for qualifying as athletic.⁸ But it

⁶The literature on the ways that sentence-meaning underdetermines what is said is vast, and the debates over what determines what a speaker says have been raging since at least the eighties. Other notable contributions include Stanley (2002, 2007), Horn (2006), and Atlas (2005). See also the papers collected in Turner and Horn (2018).

⁷By “semantic value,” here, I mean the expression’s contribution to the content of the assertion the speaker undertakes in uttering it. I do not mean what some people call the expressions “compositional semantic value.” See Ninan (2011) Rabern (2012) and Yalcin (2014) for discussion of this distinction. Notably, MacFarlane (2016, 2020a) questions even the idea that we need contextual information to determine an expression’s contribution even to assertoric content. We will touch on MacFarlane’s view of underspecification presently.

⁸See Kennedy and McNally (2005); Kennedy (2007) for discussion of gradable adjectives and vagueness, and see Sassoon (2013a,b) and D’Ambrosio and Hedden (2024) for discussion of the ways in which multidimensional adjectives depend on context.

may be the case that nothing in the context determines a specific value for any of these parameters, and I need not, and often will not, have intended a specific value for any of them. Nonetheless, my utterance may well be perfectly felicitous. This is an instance of what Jeffrey King (2018, 2021) calls *felicitous underspecification*.

Gradable adjectives are far from the only expressions that give rise to such underspecification. Consider (3):

(3) I drank all the beers. Buchanan (2010)

It is not plausible that, in uttering (3), I mean that I drank all the beers *in the world*. As we know, the domains of most, if not all natural language quantifiers are implicitly restricted by context—quantifiers are supplementives.⁹ But context will often not determine a unique restriction for the domain of the quantifier. Instead, it will leave open a range of candidate restrictions. In this case, the context will not determine whether I meant I drank all the beers *in the fridge, in the house, in the cooler, that we bought*, or any number of other restrictions. As a consequence, even relative to context, there will be a range or cloud of candidate propositional contents for my utterance.

As a third example, consider a case in which I am aiming to point out the silliness of a recent bit of political self-aggrandizement, and I utter (4):

(4) Only Ted Cruz compares himself to Galileo. King (2018)

It is standard to treat “only” as a discourse particle that introduces a set of contextually salient alternatives. But which particular alternative set has been introduced? Do I mean that Ted Cruz is the *only person in the world* who makes that comparison? Perhaps the *only person in modern politics*? Maybe the *only prominent person in US politics*? While context will surely narrow down the candidate alternative sets, it is implausible that any feature of the context comes close to determining a unique such set. Moreover, I may well have not intended any particular one of these sets, and you may well not take me to have meant a particular one—yet another instance of felicitous underspecification, in which my utterance leaves open a range or cloud of candidate propositional contents. But even these examples are just the beginning. Many, and perhaps all other supplementives—including demonstratives, indexicals, possessives, quantificational expressions, and modals

⁹For discussion of this point, see Stanley and Szabó (2000).

give rise to felicitous underspecification. The phenomenon also plausibly arises with polysemous expressions or expressions whose meaning is contested, as with political terms like “freedom,” “socialism,” “democracy,” and “fascism,” among countless others.

3. Communication and Underspecification

3.1. Underspecification in the Common-Ground Framework

When context does not fix a unique propositional content for an utterance, and instead leaves open a range of candidate contents, how does conversation proceed? In the literature there are at least three different approaches, each of which is developed within the influential framework for theorizing about conversation developed by Robert Stalnaker: the common-ground framework.¹⁰ The common-ground framework treats the context at any point in a conversation as the set of propositions interlocutors commonly accept—i.e. the common ground—and models the information in the common ground as the set of worlds consistent with the propositions in the common ground: the *context set*. Assertions, if accepted, add their content to the common ground, thus increasing the store of information shared by interlocutors and eliminating worlds from the context set.

Within this framework, the first approach to underspecification is the *diagonal* approach, endorsed by Stalnaker himself (1974/2002; 1999; 2002). On this approach, when context leaves open a range of candidate contents for an utterance, the content of the assertion undertaken with that utterance is a new proposition defined in terms of those contents: the diagonal proposition. Roughly, an utterance’s diagonal proposition is true at a world w if and only if whatever the utterance expresses at w is true at w .¹¹ On the diagonal approach, it is this proposition that gets added to the common ground when the assertion is accepted.

The second approach to underspecification, inspired by the work of Barker (2002, 2012) and developed most fully by MacFarlane (2016, 2020a,b,c),

¹⁰See Stalnaker (1974/2002, 1998, 2002, 2014).

¹¹Here I am simplifying slightly. Really, there are many propositions that meet this condition, and differ only in ways that do not affect the context set. To borrow a term from Kirk-Giannini (2018), we can call the way that an assertoric utterance updates the context set its *CS-content*. Stalnaker’s proposal is that in cases of underspecification, the *CS-content* of an assertoric utterance is the diagonal proposition restricted to the context set, the *CS-diagonal*.

is the *expressivist* approach.¹² On the expressivist view, underspecified utterances express our plans concerning how to use the terms we utter—underspecification and vagueness, for MacFarlane, are forms of practical *indecision*. MacFarlane spells out this view first by following Barker in treating the common ground as consisting of a set of world-delineation pairs. But MacFarlane treats delineations as hyperplans, where a hyperplan is a full specification of our (future) decisions concerning how to fix the extension of the underspecified expression. Given this picture of the common ground, underspecified assertions, if accepted, update the common ground by eliminating some of these world-delineation pairs, thus expressing our semantic plans.¹³

While the diagonal and expressive approaches are important, I mention them here primarily to set them aside. My goal is not to offer a general critical discussion of existing views, but to set out a view of how communication functions in the presence of uncertainty, and to examine how underspecification can be used manipulatively in its light.

To this end, the approach that serves as the main background for my own is the one developed by Jeffrey King (2018, 2021). King develops an account of updating that appeals to the idea that speakers specify only as much as necessary for the purposes or goals of the conversation. When confronted with a cloud, audiences construct a unique propositional update from the contents in the cloud by reasoning from this fact. According to King, the first step in this reasoning is as follows: insofar as the speaker left open the range of candidate contents they did, doing so must have been sufficient for their conversational purposes. From this, he claims, the audience can conclude that the speaker did not assert any specific content in the cloud, because if they had, they would have specified further. To make sense of the speaker’s underspecification, King claims, audiences must thus construct a

¹²Versions of the expressivist view have been defended, in a variety of domains, by Krifka (2012), Armstrong (2013), Coppock (2018), Kocurek et al. (2020), Charlow (2021), Mena (2023), and Soria-Ruiz (101).

¹³There is also another view, developed by Caie (2023), that is related to, but distinct from, both the diagonal and expressive views of updating. Like the expressive view, Caie treats the common ground as consisting of pairs of worlds, one of which plays a metasemantic role. But the key advance Caie’s view is that it eliminates the assumption that, even relative to a single world, an utterance expresses a unique proposition. The world that plays the “metasemantic role” may not determine a unique content. Having dropped this assumption, Caie develops a view of *disjunctive updating*, on which an assertion eliminates a pair of worlds from the context set if and only if, roughly, the first world is incompatible with the disjunction of propositions left open by the second.

propositional update built from *all* of the contents left open by the context.

According to King, this yields three candidate ways for interlocutors to update the common ground: they can update with the *disjunction* of the contents in the cloud; with the *conjunction* of those contents; or with the “doesn’t matter” update, in which all of the candidate propositions in the cloud—as well as their conjunction and disjunction—yield the same update. In order to say how interlocutors choose between these three updates, King appeals to the idea that conversation is structured around a *question under discussion* (QUD), which is often integrated into the common-ground framework to help model the topic and goals of conversation.¹⁴ The QUD is typically treated as partitioning the context set into mutually exclusive and exhaustive sets of possibilities, each of which represents a complete answer to the question at hand. The goal of conversation is then understood as follows:

Answer the Question The goal of conversation is to make a complete answer to the QUD common ground.

Roberts (1996); Stokke (2018); King (2021)

This conception of the goal of conversation generates norms or principles that govern how interlocutors contribute to conversation. In particular, the QUD framework implements the Gricean principle of Cooperativity as follows:

QUD Cooperativity When the QUD is Q , assert the strongest partial answer to the QUD that you know.¹⁵

Thus, in the QUD framework, interlocutors all have the goal of answering the QUD, and the fact that each interlocutor cooperatively pursues this goal is part of the common ground by default. Thus, the QUD framework treats conversation as a *cooperative* game of *complete information*: it is a game in which the goals of interlocutors are aligned and known to align on providing an answer to the QUD. Given this, in constructing their update, audiences can rely on the fact that the speaker is aiming to offer the best answer to the QUD that they can, while also specifying no more than necessary to achieve their goals. Accordingly, King claims that audiences update with

¹⁴Roberts (1996) is the canonical presentation.

¹⁵Different authors define QUD cooperativity in different ways. For instance, Dorst and Mandelkern (2023) define their norm for assertion in terms of certainty, though this is arguably too strong to serve as a plausible norm governing assertion.

the weakest candidate propositional update such that no stronger candidate update provides a better answer to the QUD—in short, they update with the disjunction unless the conjunction provides a better answer to the QUD.¹⁶

While King is surely correct that in cases of underspecification, audiences decide how to update by reasoning about the speaker and their goals, King’s view nonetheless faces two serious problems. The first problem stems from the fact that felicitous underspecification can arise even when the speaker is intending to communicate some specific content. Communication is often uncertain, and speakers often rely on audiences to infer or guess the contents of their assertions under uncertainty, on the basis of the evidence in the context. Moreover, underspecifying in this way is often rational—specifying fully comes with additional messaging costs, which the speaker may not be willing to pay if they are sufficiently confident that the audience will recover their intended meaning. It follows from this that in cases of underspecification, context will leave also open whether the speaker intends to communicate some specific content in the cloud or whether they had no specific content in mind. In such cases, audiences will not be licensed to conclude that the speaker has asserted either the disjunction or the conjunction of the contents in the cloud, as King claims. Rather, audiences will be forced to reason under uncertainty about what exactly the speaker has asserted and how specific their communicative intentions are.

We can illustrate this point by again considering the case of **Electoral Strategy**. To simplify slightly, let’s suppose that you ask me “what are you going to do about gun crime?” and I respond “Everything in my power.” My response is surely felicitous, even though it underspecifies its content and context leaves open a range of candidate contents. But it seems clear that, for all you know given the information in the context, I may have meant something specific—one of the particular contents left open by the cloud—or something general—i.e., some disjunction or conjunction thereof. In ruling out the former possibility, King’s view makes incorrect predictions about how the audience should update when they encounter a cloud.¹⁷ For suppose

¹⁶Stokke (2018) offers a similar, question-sensitive proposal about how to recover a unique content. First, Stokke introduces a notion of minimal content, and claims that in underspecifying, what a speaker says is the weakest answer to the QUD that entails or is entailed by this minimal content.

¹⁷This is not just true of King’s view, but plausibly true of any view motivated by the principle of UNIFORMITY, which tells us that in order for rational communication to occur—i.e., in order for an assertion to update the common ground—each interlocutor must be certain of its content.

that I intend to communicate something specific with my utterance—that I plan do everything in my power *including challenging the second amendment*—and that context, despite leaving open a cloud, provides you with very good evidence that this is what I intend. In this case you may, on good grounds, take me to have meant just this and update your presuppositions accordingly.¹⁸ In this case, it is clearly wrong to think that you must update with the conjunction or disjunction. Rather, I have asserted something specific and you’ve correctly determined what I have asserted under uncertainty, and you’ve changed your presuppositions to reflect that.¹⁹ I take it that communicative interactions of this form—in which audiences do their best to infer what the speaker has asserted under uncertainty—are ubiquitous. As we will see presently, it is the ubiquity of these situations that makes manipulative underspecification such a powerful strategy.²⁰

The second problem for King’s view arises as a consequence of his appeal to the idea that, in selecting an update, audiences know, and can reason from the premise that, the speaker is abiding by QUD cooperativity. Given this premise, King’s update rule yields that interlocutors update with the weakest candidate update, i.e., the disjunction, unless some stronger content—the conjunction—provides a better answer to the QUD. But in strategic contexts, such as **Electoral Strategy**, speakers often have goals that conflict with answering the QUD; they have *non-communicative* goals, such as getting likes, manufacturing perceived agreement, saving face, avoiding conflict, minimizing commitments, or fostering solidarity. Guided by these goals, speakers may deliberately specify less than needed to answer the QUD. Moreover, in impoverished contexts, audiences will typically be

¹⁸As we will see more fully below, updating in this way will involve some risk, because you are not certain that this is what I’ve said. But you nonetheless have very good evidence that it is, and it seems clear that interlocutors do often come to presuppose or take things for granted in conversation—i.e., treat them as common ground—without being certain of them, given the risks and rewards of doing so.

¹⁹For a similar argument against the diagonal approach to assertoric content, see Kirk-Giannini (2024). Kirk-Giannini’s argument functions by considering a case where it is common ground that the speaker intended something specific, but not common ground what it is. In this case, he argues, the content of the assertion is not the diagonal proposition.

²⁰Going forward, I will sometimes say that audiences “guess” what the speaker has asserted, but I use this term as a non-technical gloss for the candidate contents the audience can take the speaker to have asserted, and so the candidate ways they can update their presuppositions if they accept that assertion. In other work, I develop a view of updating under uncertainty that draws on the work of Holguin (2022) and Dorst and Mandelkern (2023) which takes the idea of guessing more seriously.

uncertain about the speaker’s communicative goals, and so be uncertain about whether they should in fact update with the proposition that best answers the QUD. Any account of updating that appeals to conversational goals must take into account that speakers sometimes underspecify strategically, and audiences, in trying to recover the contents of such assertions, must reason under uncertainty about whether the speaker is being fully cooperative.

3.2. A Game-Theoretic Approach

What we need is an account of how an audience updates under uncertainty that is sensitive to their credences about what the speaker has asserted, including whether the speaker had something specific in mind, as well as their credences about the speaker’s conversational goals. Here is a sketch of just such a view. Let’s suppose that in a context, an utterance u leaves open a range or cloud of fully specific contents R . We can capture the audience’s uncertainty about the content of a speaker’s speech act using a conditional probability distribution over R , $P(R|U)$: i.e. the probability that the speaker meant one of those specific contents conditional on their having uttered u . The contours of this distribution will be partially fixed by the audience’s priors, which will be determined by their presuppositions at the time of the utterance.

Given this distribution of credences over candidate contents, we then need to ask how the audience, upon hearing u , selects an interpretation, updates their presuppositions, and chooses subsequent conversational actions. The approach I will adopt here is broadly game-theoretic, and treats the audience as selecting a content from the distribution in a way that maximizes their expected utility, as in so-called “ambiguity” or “interpretation” games.²¹ The key idea of such games is that, given the audience’s distribution of credences over the particular contents left open by the context—i.e.,

²¹For an extended discussion of games of ambiguity, see Parikh (1991, 2001, 2010). For discussion of interpretation games, see Asher and Lascarides (2013). The game-theoretic approach is only one of many options for selecting a content under uncertainty. We could, for instance, appeal to another form of inductive inference, such as maximum *a posteriori* (MAP) estimation or inference to the best explanation. Another alternative is to bypass updating altogether, go fully Bayesian, and treat addressees as simply having a distribution of credences over candidate contents, and—without coming to any determination about what was asserted—use their utilities to decide on subsequent conversational actions. But here my aim is to construct a model that integrates uncertainty into the framework in which interlocutors update their presuppositions, so I treat speakers as selecting a content or interpretation from the range of options.

their views on what the speaker has asserted—they must decide which candidate content the speaker is proposing to add to the common ground, given the risks and rewards of the various options, which are represented by their utilities.

As we have seen, however, in undertaking an assertion, speakers need not, and often do not, have a fully specific proposition in mind. Accordingly, we also need to allow for the possibility that the audience may not take the speaker to have asserted anything fully specific at all. The most straightforward way to do this is to expand the range of candidate contents over which the audience has credences to include disjunctive (and conjunctive) contents. The audience can then assess, in light of the information in the context at the time of the utterance—including their beliefs about the speaker’s goals—how specific they take the content of the speaker’s assertion utterance to be. Here I will simplify by assuming that, in addition to the particular, fully specific contents left open by the context, there is just one disjunctive content in this range: the disjunction of all contents left open by the context at the time of the utterance.²²

To illustrate, consider the following case. Suppose we’re at a dinner party. It is common ground that I either want to leave to get ice cream, leave to go to another party, or leave to go home, but you are uncertain which. It is also common ground that I often like to get ice cream after dinner. I utter (5):

(5) I’m ready.

Here, context leaves open three candidate fully specific contents for my utterance:

1. C_1 : I’m ready to go get ice cream.
2. C_2 : I’m ready to go to another party.
3. C_3 : I’m ready to go home.

But context also leaves open the possibility that in uttering (5), I didn’t have any of these specific contents in mind, and in fact meant the disjunction $C_1 \vee C_2 \vee C_3$. Suppose that, given the information in the context, your credences over these four possibilities are as follows:

²²The “best-guess” model of updating (D’Ambrosio, manuscript) is a closely related alternative approach to dealing with nonspecific updates. On that view, audiences can update with more and less specific contents depending on their background views of how specific the speaker needs to be to answer the QUD.

- $P(C_1|U) = .5$
- $P(C_2|U) = .1$
- $P(C_3|U) = .1$
- $P(C_1 \vee C_2 \vee C_3|U) = .3$

You then must consider the various risks and rewards of taking me to have asserted each of these four candidate contents. If you take me to have asserted C_1 , update your presuppositions accordingly, but are wrong, there will be a small cost, since we've miscommunicated. But if you are right, there will be a benefit, because you've correctly inferred or guessed what I have asserted under uncertainty. Similarly, if you take me to have meant the disjunction, and do not take for granted that I meant anything specific, the potential costs may be lower—you did not take me to have meant something specific—but so may the potential benefits—we've failed to exchange as much information. Thus, as we will see in more detail in §7, your utilities represent the conversational costs or benefits of taking me to have asserted one or the other of the contents left open by the context, given what I have actually asserted.

The resulting view solves the two problems for King's view raised above. First, the game-theoretic model allows for audiences to update their presuppositions with a specific content even when there are other such contents not ruled out by the common ground. To speak colloquially, it allows them to update with a *guess* about the content of the speaker's assertion. Such updating is risky, yes, but will often be worth it, given the potential rewards of successfully communicating. But the view also allows the audience to update with the disjunctive content if they are risk-averse, and want to be sure not to take for granted something more specific than the speaker intended. Thus, on this view, King's disjunctive update is just one of wider range of updates—including each of the candidate specific contents—available to the audience when they encounter an underspecified utterance.

Second, this view does not assume that the speaker is being fully cooperative and aiming to specify as much as necessary to answer the QUD. Rather, on the game-theoretic view of communication, speakers may have intentions of varying degrees of specificity, and audiences must reason under uncertainty about what the speaker has asserted and what their conversational goals are. As we will see more fully below, these two forms of reasoning interact. How the audience updates will be sensitive to their background beliefs about the content of the speaker's assertion and their beliefs about the

speaker's goals, and so will take into account the possibility that the speaker is not being fully cooperative. If the audience thinks that the speaker has deliberately underspecified, they will have the option to treat the speaker as having uttered something infelicitous.

4. Strategic Underspecification

4.1. Examples of Strategic Underspecification

With this model of updating as background, let us now consider some further examples of how underspecification can be deployed strategically.²³ Consider **Vaccine Evidence**:

Vaccine Evidence I'm a public health administrator, and I'm trying to get as many people as possible to take a vaccine against a recently discovered and highly dangerous disease. There are countless studies that have been run concerning the vaccine, and most of the data indicates that the vaccine is safe and effective, although some small studies point toward potentially serious side-effects. On the one hand, I don't want you to take me to have lied or exaggerated the evidence for their safety, but I also don't want to admit something that will keep you from getting one.

Given my goals and my uncertainty, a wise choice is to utter the following:

(6) There's no compelling evidence that the vaccine is dangerous.

Here, context does not plausibly fix a unique restriction for the quantifier "no compelling evidence," nor does it fix a unique threshold for what makes some evidence "compelling." Instead, context leaves open a range of candidate domain restrictions and thresholds: I might mean no evidence *whatsoever*; I might mean no evidence *that we've collected*; I might mean no evidence collected *by reputable and well-run studies*; or I might mean no evidence *that would warrant not getting vaccinated*, to name just a few. Importantly, some candidate fully precise meanings will exclude the small studies from

²³Some of the ways that underspecification can be used strategically have been examined in the literature in political science (Aragonès and Neeman, 2000; Aragonès and Postlewaite, 2002; Tomz and van Houweling, 2009), formal pragmatics (Jaegher, 2003; van Deemter, 2010; Blume and Board, 2014), and organisational communication (Eisenberg, 1984).

the domain of quantification, while others will not. In this situation, you will be forced to take a guess—guided by your credences—about what I have asserted. You may guess that I have asserted something specific—one of the specific contents left open by the context—you may guess that I have asserted something general—for instance, the disjunction of these contents—or you may guess that I have strategically underspecified, and so treat my speech act as infelicitous.

Consider another example:

Troop Withdrawal You're a high-ranking general and I'm a low-level analyst at the Pentagon. You're known for having a temper, don't like people contradicting you, and tend to belittle people who do. I'm in the uncomfortable position of being grilled by you in a meeting without having any idea about your views on the issue at hand. You ask: where should we station our troops?

I reply:

(7) We have to station them in the safest available location.

In (7), “have to” is a modal, and as we know, fixing the domain of a modal requires supplementation from context. Likewise, “safest available” is an adjectival phrase that is both vague and multidimensional, and its semantic value will depend heavily on context. But again, since the context is impoverished, it will not single out a unique semantic value for either supplementary, and so will leave open a range or cloud of candidate contents for my utterance. As in the previous case, using language of this kind forces you to reason under uncertainty about what I mean and whether I'm speaking strategically. You will need to guess, based on the evidence provided by the context, what exactly I mean, how specific my intentions are, and whether I'm speaking strategically, and then decide what to take for granted as conversation proceeds.

Now consider **Cluster Bombs**:

Cluster Bombs: I'm a politician deciding whether or not to send a powerful but controversial weapon to an ally to aid them in a fight against an invading power. In a televised interview, you ask me to justify my decision to send the weapon to our ally, given that many other countries have condemned the decision and said that there were other ways to proceed. You ask: weren't there other equally good ways of supporting our ally?

I respond:

- (8) Sending the controversial ammunition to help our allies was a difficult decision, but it was necessary.

Here again, necessary *for what?* Clearly, sending them the weapons was not *metaphysically* necessary. What kind of necessity am I invoking? And how is the modal base fixed? Do I mean necessary *for them to win the war?* *Necessary to achieve our country's strategic goals?* *Necessary to preserve our alliance?* Here, again, context will not settle the matter. Instead, you will be forced to guess what I mean. For all you know, I may have meant something specific, I may have meant something extremely general, or I may have been deliberately underspecifying so as to allow myself latitude in downstream conversation.

4.2. Advantages of Underspecifying

Underspecification provides me, the strategic speaker, with a range of advantages, and it is worth cataloguing these advantages explicitly. The first advantage, which we've already seen, is a consequence of the uncertainty occasioned by underspecification:

Keeps you Guessing: Underspecification forces an audience to come to a determination about the content of an assertion under uncertainty and choose subsequent conversational actions on this basis.

In other words, underspecification forces the audience to take a guess about what act the speaker is undertaking. One consequence of the fact that you must guess at what I assert is that it allows for the possibility that you will *read in* a particular content, but one that I wasn't in a position to assert outright. For example, suppose that in the case above, because you are confident that I am aiming to say something both true and informative, you have credence .9 that the content of my assertion is that sending the ammunition is necessary *to preserve our alliance*. It seems perfectly rational for you to conclude that I have asserted that thing, and update your presuppositions accordingly. But in taking me to have asserted something specific while I have deliberately underspecified—i.e. in virtue of reading in—you risk being deceived about what I've asserted in a way that is conducive to my goals and works against your own.

The second advantage is that underspecification makes it difficult for you to reject my speech act:

Makes Rejection Difficult: In leaving an audience uncertain about what was meant, underspecification makes rejecting or challenging a speaker’s speech act difficult.

When strategic underspecification is done well, at least some of the contents left open by the context will be ones that the audience cannot or would not reject, and instead would accept. But when some such contents are among the candidate contents left open by the context, it will be difficult for the audience to reject the speaker’s speech act. One reason is that, if the audience takes the speaker to be cooperative, then they will be more confident that contents they are willing to accept are the content that has been asserted—there will be something like a principle of charity in how they select an interpretation under uncertainty.²⁴ But this entails that, under uncertainty, acceptable contents are more likely to be the ones that audiences take to be the assertion’s content.

In cases of underspecification, the only thing to which you can be certain I’ve committed is something so vacuous that you cannot reasonably reject it. This interpretation—the weakest interpretation, i.e. the disjunction of all the candidate specific meanings—provides a lower-bound on my commitments. In order to reject my assertion, you will need to guess that I mean something more specific than this disjunction—you will need to read in something that you reject. This will happen most often in the case where your priors make you very skeptical of my motives. Only then will you be in a position to guess that I mean something specific which you can reject. Of course, you may also ask for clarification in order to remedy your uncertainty. But clarification requests also have costs. As Deigan (2022) shows, clarification requests can be costly because they take time, express ignorance, and often appear to challenge authority.²⁵

However, failing to reject my assertion while uncertain about its content leaves you in a vulnerable position due to the third and fourth advantages

²⁴This advantage depends on how strong we take the presumption of cooperativity to be. The speaker may be cooperative but incompetent, and so we may not interpret them charitably. But such cases aside, it is plausible that interpretive charity is often at work in fully cooperative conversation.

²⁵Deigan (2022) discusses the costs of requests for clarification in his discussion of *stupefying*: the phenomenon in which an audience accepts an assertion without understanding it. There are important similarities between stupefying and deliberate underspecification. In fact, it is plausible that stupefying is a form of deliberate *overspecification*. In both cases, it is difficult for an audience to reject the speaker’s speech act because they do not understand it, and instead are at best uncertain about its content.

of underspecification. The third advantage is that underspecification allows me to gauge your reaction and calibrate subsequent utterances in light of it:

Allows Gauging and Calibration: Underspecification allows speakers to draw out information about their interlocutors and so reduce their uncertainty about what to say in downstream conversation.

The fourth advantage of underspecification is that it gives me the prerogative to retroactively change my commitments in light of the information elicited from you:

Allows Changes in Commitment: Underspecification allows for strategic retroactive changes in commitment—both plausible deniability and plausible assertability.

cf. Peet (2016, 2017); Weiser (1974)

Together, these two advantages are perhaps the most powerful and insidious elements of underspecification. As conversation progresses and new information comes to light, I can retroactively commit to whatever interpretation I now take to be most advantageous, and pretend or insist that I meant that thing all along. Of course, such retroactive commitments may be resisted. However, since you were not in a position to rule out that I meant that thing all along, such retroactive changes in commitment will typically be plausible, and resistance will be difficult.²⁶ Further, you will be aiming to resist such changes *after* having accepted my initial speech act, which makes such challenges less likely to be successful.²⁷

To see this dynamic in action, consider again my felicitous but underspecified utterance from **Electoral Strategy**, repeated here:

(9) **Politician:** We should do *everything in our power* to keep guns out of the hands of *dangerous* people.

Consider one way that the conversation might continue:

²⁶It is an open question how this view of what makes deniability plausible interacts with views such as those developed by Dinges and Zakkou (2023) and Peet (2016). In the interests of space, I will here set discussion of this question aside.

²⁷A reviewer helpfully points out that the advantages of manipulative underspecification are similar to the advantages of dogwhistles. I completely agree, and in fact think that manipulative underspecification may well be one phenomenon that gets called “dogwhistling”. But importantly, manipulative underspecification concerns the content of an assertion itself, rather than inferences audiences can make based on the content of an assertion. See Khoo (2017), and Saul (2018, 2024) for discussion of dogwhistles.

- (10) a. **Constituent:** Great, but what are you going to do about all of the guns already in circulation?
- b. **Politician:** Requiring background checks for all gun purchases is just a first step—eventually we hope to implement a comprehensive buyback program.

In your response, you (the constituent) first accept my (the politician’s) initial speech act—it would be hard not to—and then reveal something about your views: that you are, broadly speaking, anti-gun. This allows me to retroactively specify and commit myself to having all along intended something to which I now know you are amenable: instituting required background checks. Thus, in virtue of having left you uncertain about the content of my initial assertion, I can plausibly and retroactively commit to one of the specific contents left open by the initial context.

But even if you do manage to challenge my assertion, the flexibility afforded by underspecification still allows me to achieve my strategic aims. Consider an alternative continuation:

- (11) a. **Constituent:** But what about the second amendment? Law abiding citizens have the right to bear arms!
- b. **Politician:** Those are exactly the people I’m going to make sure are allowed to keep their guns.

Here you challenge my assertion, thinking that I’ve committed myself to a broadly anti-gun policy. But in challenging me, you again give me information about your views, which, in virtue of having underspecified, I can exploit: again I can retroactively, and plausibly, insist that I meant something determinate and more amenable to your views all along.

Now consider one way **Vaccine Evidence** might continue after I utter (12):

- (12) There’s no compelling evidence that the vaccine is dangerous.
- (13) a. I know, they’re just sowing doubt, and people are dying because of them!
- b. They don’t care that there isn’t a single shred of evidence for their claims.

Here, you agree with my initial claim—again it would be hard not to—and respond in a way that reveals your views. I then commit to something stronger that aligns with your views—I commit to one of the stronger specific contents initially left open by the context. But now consider an alternative

continuation, in which you offer a more measured response:

- (14) a. Well, there is some limited evidence that there can be bad side-effects, but it's almost certainly outweighed by the benefits.
- b. Of course that's what I meant—there simply isn't any evidence that's good enough to warrant not getting vaccinated.

Here you recognize that my claim is only plausible when the quantifier is restricted, and so you go on to challenge the stronger version of my claim, where the quantifier is unrestricted. I then retreat to the meaning in the cloud on which the quantifier “no evidence” is restricted by “that's good enough to warrant not getting vaccinated.” Thus, I plausibly deny having asserted the stronger content, and commit retroactively to having meant the weaker one.

Retroactive commitment or specification resembles the conversational strategy of adding contents to the common ground by presupposing them, or making them not-at-issue in a particular speech act. This strategy is sometimes called making use of the “conversational back door.”²⁸ The conversational back door serves as a way for a speaker to obviate an audience's defenses against adding a content to the common ground. Presupposing rather than asserting a content makes it difficult for an audience to keep that content from entering the common ground, at least partly because they may not fully attend to that content, and because conversational dynamics make it difficult to challenge content that is not-at-issue. Similarly, underspecification and retroactive commitment make it difficult to challenge an initial assertion, because the audience is uncertain about what was asserted in the first place, and it is only after this acceptance that their uncertainty is resolved, and the speaker's contribution to the common ground is settled. Thus, one might say that underspecification and retroactive commitment provide speakers with a conversational “side door” in addition to a back door.

5. Manipulative Speech as Covertly Strategic Speech

We now have an outline of how strategic underspecification works in conversation. In order to realize some non-communicative goal, such as manufacturing perceived agreement, gaining status, or avoiding conflict, or sneaking a content into the common ground, a speaker underspecifies the

²⁸For discussion of back-door speech acts, see Stanley (2015) and Langton (2018).

content of their speech act, leaving open a cloud or range of candidate contents. The audience guesses or selects an initial interpretation based on their presuppositions at the time of the utterance. But once the audience chooses this interpretation and responds, the speaker can attempt to retroactively commit to an alternative interpretation left open by the context, and pretend or insist that they meant that all along; they aim to make their underspecification covert. In aiming to get the audience to believe that they have all along committed to one or another determinate content, the speaker's speech is not only strategic, but *manipulative*. Manipulative speech is speech that is *covertly strategic*.

Recall that above, a strategic context was defined as one in which my primary goal, as speaker, is not to communicate with you, where communication is understood as sharing knowledge. Strategic speech is speech that aims to realize a goal that conflicts with sharing knowledge; strategic speech always occurs in a strategic context.²⁹ If we in turn understand the goal of sharing knowledge as the goal of answering the QUD, then we can understand strategic speech as speech that deliberately violates QUD Cooperativity in service of some non-communicative goal. In other words, strategic speakers have reason not to provide their interlocutors with a complete answer to the question under discussion.

Strategic speech thus involves a speaker deliberately choosing not to do what is optimal in order to share knowledge with their interlocutors. This brings out the connection between strategic speech, Grice's cooperative principle, and the conversational maxims. One natural way of understanding the Gricean maxims—the way which Grice himself seems to understand them—is as specifying the optimal way to exchange information or share knowledge. Understood this way, we can see the maxims as providing a rough guide to the varieties of strategic speech—any way of violating a maxim in order to realize a non-communicative goal is a form of strategic speech. Strategic speakers may lie, provide too much or too little information, be deliberately unclear, say things that are irrelevant, presuppose rather than assert something crucial, deliberately underspecify, or employ any number of other strategies, each corresponding to a different way of deliberately violating a maxim and so failing to be fully cooperative.

²⁹Of course, the idea that communication consists in sharing knowledge or information is an idealization, one made by Grice (1989), Lewis (2002), Stalnaker (2014), and many others. But the definition of strategic speech can be generalized from the goal of sharing knowledge to any goal that conversation might have. The key point is that the speaker pursues a goal that conflicts with the audience's goal for the conversation.

Strategic speech can be useful when it is completely overt—i.e., when it is obvious that a speaker is not being fully cooperative. For instance, in cases of bald-faced lying, the speaker’s goals are known, and yet lying is still advantageous to the speaker. Similarly, a speaker may refuse to answer a question by changing the subject, this may be obvious to their audience, and yet doing so may be advantageous. Nonetheless, strategic speech will often be most useful when it is not recognized as strategic by the audience—i.e., when it is *covertly* strategic. When a speech act is covertly strategic, the speaker poses as cooperative—poses as aiming to share knowledge or provide a complete answer to the QUD—when in fact they are pursuing an ulterior, non-communicative goal. This covertness, on my view, is what distinguishes manipulative speech from merely strategic speech.³⁰ Manipulative speakers pose as cooperative and so hide their non-communicative goals. Manipulative speakers are, so to speak, *false friends*.

Covertness, however, is not an all-or-nothing affair. Contexts such as the ones in our examples above are what I called *impoverished* contexts. In an impoverished context, the information shared by interlocutors—i.e., the common ground—leaves the audience *uncertain* about the speaker’s goals and so uncertain about whether they are speaking strategically. In such contexts, the audience must reason under uncertainty about the speaker’s goals in order to determine what actions to take in subsequent conversation. To take the simplest example, in ordinary cases of lying, audiences reason under uncertainty about whether the speaker is deliberately asserting a falsehood. Whether to accept the speaker’s assertion or reject it depends on this reasoning. The audience’s credence that the speaker is not speaking strategically, and that their primary goal is communicative, is what I call their degree of *linguistic trust*. When an audience has a high degree of linguistic trust, they are more vulnerable to being manipulated, for linguistic trust is exactly what manipulative speakers exploit to realize their non-communicative goals.

Manipulative underspecification is an instance of this more general category. In manipulatively underspecifying, a speaker leaves open a range of candidate contents in a way that is not fully cooperative, and is not optimal for answering the QUD. The audience, uncertain about the content of the assertion, has to determine what has been asserted by reasoning under uncertainty about the speaker and their goals. When the audience is trusting,

³⁰The discussion in this section draws on the account of strategic and manipulative speech I develop in D’Ambrosio (forthcoming).

and takes the speaker to be fully cooperative, this will affect their interpretation of the utterance—it will lead them to guess that the speaker has said something optimal for answering the QUD. In doing so, they may read in a content that is advantageous to the speaker, and helps the speaker realize their ulterior, non-communicative goal.

The reasoning that leads to such reading in is similar to the reasoning that occurs in calculating a strengthening implicature.³¹ When a speaker obviously fails to give as informative an answer to the QUD as they can—for instance, by leaving out some crucial information—the audience, presuming that the speaker is cooperative, can infer that they must have meant something stronger than what they said. If, for instance, I say that Mary and I went to the store, on the basis of the presumption of cooperativity, you can infer that we went *together*. Manipulative underspecification involves similar reasoning, except it concerns the content of an assertion. When a speaker strategically underspecifies, instead thinking of the speaker as having asserted the disjunction of contents left open by the common ground, the audience, presuming the speaker is fully cooperative, strengthens the speaker’s assertion, taking them to have said something more specific and conducive to answering the QUD than the mere disjunction.

However, manipulative underspecification, like manipulative speech more generally, need not always work against the interests of the audience.³² On the contrary, such underspecification may be mutually beneficial when interlocutors are uncertain about one another’s goals and background beliefs. Consider, for instance, a situation where two interlocutors are trying to assess mutual romantic interest under uncertainty. In such a context, the speaker may deliberately underspecify, hoping to draw out further information about their interlocutor’s interest or feelings. If the gambit is successful, the speaker will come to know more about their interlocutor, and can then retroactively commit to one or another more specific reading—whichever is most advantageous, given their new information—and aim to make it common ground that they intended that content all along. Such retroactive commitment and update will often be in the interest of the audience as well, allowing them to proceed in conversation as if no advance has been made.

³¹For discussion of strengthening implicatures, see Horn (1984), Davis (2019).

³²For discussion of how manipulation and manipulative speech may serve the interests of the audience, as in cases of paternalistic manipulation, see D’Ambrosio (forthcoming).

6. Pied Piping and Contextual Fragmentation

What kind of speech act, if any, is undertaken by a speaker who manipulatively underspecifies? There are several reasons to think that in deliberately underspecifying the content of an assertoric utterance, a speaker does not make an assertion, and indeed, does not engage in a genuine speech act at all. The first reason is that the speaker does not have a Gricean reflexive intention to communicate any particular proposition. In the situations above, I don't know what your views are, and so don't know what proposition to communicate in order to achieve my goal—I don't even know enough to lie. But if I don't know what proposition I need to communicate in order to achieve my goal, I do not intend to communicate that proposition. And while one could say that I intend to communicate whatever content will achieve my goal, I do not want to be recognized as so intending. Rather, which exact speech act I want to be understood as having undertaken will depend on facts that are unknown to me when I undertake the utterance, and which will typically only emerge in downstream conversation.

However, even if I do not have a Gricean reflexive intention to communicate any particular proposition, a natural thought is that the content of my assertion is something weak or general, like the disjunction of the propositions in the range R left open by the context. King's view often delivers this verdict.³³ However, on his view, in order to assert this disjunction, the speaker must have an intention to be recognized as associating their utterance with a particular range R of fully specific utterances. But in the strategic cases just discussed, I do not have such an intention. On the contrary, I do not wish to be recognized as leaving open a range of candidate propositions. If something specific would be most advantageous, then I want to be taken to have meant that all along. Thus, there is neither a range of propositions that I intend the audience to recognize I have left open, nor is there any particular proposition that I want the audience to recognize as the one I intend to communicate. Thus I do not have a Gricean reflexive intention to communicate any proposition whatsoever.

A final reason to think that instances of manipulative underspecification are not genuine assertions comes from their effects—or lack of effects—on the common ground. Recall that on the standard, Stalnakerian view of assertion, the characteristic effect of an assertion is to add its content to the common ground, and so contribute to interlocutors' store of shared informa-

³³As do Stokke (2018) and Caie's (2023) views.

tion. Typically, instances of manipulative underspecification will not affect the common ground *at all*. The reason is that, frequently, how an audience updates their presuppositions—i.e. which act they take the speaker to have engaged in, given their background beliefs and evidence—will not itself be transparent to the speaker, at least not at first. In impoverished contexts, just as the audience will be guessing about the content of the speaker’s assertion, the speaker will be guessing about how audience updates their presuppositions concerning what has been asserted. When successful, deliberately underspecified assertoric utterances affect the context by *deceiving* the audience about what speech has been undertaken, and giving them false beliefs about the common ground, but in a way that is initially opaque to the speaker. Thus, such utterances add nothing to the common ground.

Let’s say that when a speaker deliberately and manipulatively underspecifies the content of their speech act, but also intends to be taken to have meant something fully determinate and conducive to their goals, the speaker is *pied piping*. Pied piping is what Bach and Harnish (1979, p. 101) call a *covert collateral act*, a category which includes “acts of manipulation, including such devious acts as innuendo, deliberate ambiguity, and ‘sneaky presupposition’.” Such acts, Bach and Harnish argue, are mock speech acts—they do not have illocutionary force, because the intentions with which they are undertaken are not reflexive, and instead must remain hidden for the act to bring about its perlocutionary aim.

Within this category, pied piping is best understood on the model of deliberate ambiguity, which Bach and Harnish characterise as follows:

In deliberate ambiguity clearly the speaker does not have an *R*-intention [a reflexive intention] that his utterance be taken in one of two ways. Rather, since his deliberate ambiguity is covert, he has a simple intention that the hearer take him to have one of two possible *R*-intentions, which one to be determined by the subsequent course of conversation.

Bach and Harnish (1979, p. 102)

Drawing on the work of Ann Weiser, Bach and Harnish go on to make the following claim:

Assuming the hearer does not recognize the speaker’s intention to be ambiguous, the speaker has direct control over how the utterance will be taken, given the response.

ibid., citing Weiser (1974)

This captures the nature of underspecification and its conversational dynamics perfectly. In pied piping, a speaker succeeds in realizing their goal—that of appearing to agree with someone, of being liked, of avoiding conflict, or of gaining approval—just in case they are not recognized as having deliberately underspecified, or at least if they can plausibly insist that they have not. Moreover, insofar as the speaker is not recognised as having deliberately underspecified, they have the prerogative to select which content is added to the common ground from among those left open. This gives us yet more reason to think that rather than being a genuine speech act, pied piping is a non-illocutionary, mock speech act that, when successful, is taken to be a genuine illocutionary act with a determinate content.³⁴

While pied piping does not have a characteristic effect on the common ground, it does have a characteristic effect on the context: it *fragments* the context. We can think of contexts as individuated by the presuppositions of interlocutors, shared or otherwise. The common ground is constituted by the set of these presuppositions that interlocutors share. But when pied piping is successful, the audience comes to presuppose that an assertion with a particular content has been made, when in fact no such speech act has occurred. Further, unless the speaker knows how their interlocutors guess, they will be uncertain about how the audience has updated their presuppositions—the audience’s updates will be opaque to her. Pied piping thus has the effect of getting the audience to presuppose something that the speaker does not—it has the characteristic effect of creating a *defective context*. The deliberate creation of a defective context by means of underspecification is what I call *fragmentation*.

However, as we saw above, the context will often not be defective for long. If all goes well for the speaker, in the course of the conversation, the audience will reveal which of the various interpretations they adopted, and perhaps also which of those interpretations was the most advantageous one for the speaker. In revealing which interpretation they adopted, the audience makes it common ground what they took the speaker to have asserted. This

³⁴Are there forms of manipulative underspecification other than pied piping? I wish to leave this question open. I take Bach and Harnish’s characterization of the speaker’s intentions in deliberate ambiguity, suitably generalized, to be necessary and sufficient for pied piping to occur. One candidate for being an instance of manipulative underspecification that does not satisfy this condition is the case in which a speaker has a particular *R*-intention to communicate a proposition, but underspecifies merely to preserve plausible deniability. The question is whether this form of underspecification is manipulative. This is not a question I will try to settle here.

is a process of *defragmentation*: it once again unifies the information-states of interlocutors, giving them (at least approximately) aligned views of what has taken place in the conversation. If it turns out that the audience has read in an interpretation that is advantageous to the speaker, the speaker can then go on and retroactively commit to having undertaken that speech act all along. If the audience has read in a content that is not advantageous to the speaker, the speaker can deny having undertaken that speech act, and insist that they undertook one with a more favorable alternative content.

7. A Game-Theoretic Model of Contextual Fragmentation

Standard models of conversation in pragmatics, including the common-ground model discussed above, treat conversation as a *cooperative game of complete information*—interlocutors cooperate, and are known to cooperate, in working toward a shared conversational goal. But in cases of manipulative underspecification, a speaker underspecifies strategically, leaving interlocutors uncertain about the content of their speech act and about their goals in undertaking it. An adequate model of manipulative underspecification must capture how the goals of speaker and audience conflict, how interlocutors reason under uncertainty about the speaker’s goals and the content of their utterance, and how their subsequent conversational actions depend on this reasoning.

In game-theoretic pragmatics, the standard way of modelling linguistic exchanges in which interlocutors have incomplete information about each other’s goals is with a particular kind of Bayesian game called a *signalling game*.³⁵ Each signalling game involves a set of interlocutors—here we will

³⁵See Franke (2009, 2010, 2013), Franke and Jäger (2016), Parikh (1991, 2001), among others. Signalling games are also used in the now-well-known Rational Speech Act (RSA) framework for pragmatics (See Frank and Goodman, 2012; Goodman and Frank, 2016; Lassiter and Goodman, 2017). However, the RSA model uses these games very differently than I will here. Crucially, the rational speech act model does not, strictly speaking, involve multiple agents. It aims to model pragmatic reasoning from the perspective of a sophisticated listener, who draws conclusions by reasoning about what a speaker should utter, given their goals. Second, the RSA model presumes that linguistic expressions have fixed interpretations to which the listener has access in pragmatic reasoning. It does not address uncertainty about the contents of assertions themselves. Third, the listener, in the RSA model, does not have utilities—interpretation is purely epistemic. Fourth, and relatedly, the RSA model does not involve updates to the common ground. The model I offer here aims to integrate signalling games with the idea that speech acts update the common ground.

assume that there are just two: speaker and addressee. In addition, a signalling game has three core components: a set of speaker types, a set of messages, and a set of addressee actions, where what is in these sets is commonly known. Further, speaker and addressee both have utility functions representing their goals, and they have probabilistic beliefs about facts that are not common knowledge. The speaker has probabilistic beliefs about what action the receiver will undertake upon receiving certain messages, and the receiver has probabilistic beliefs about the speaker’s type and what message the speaker will send given their type.

When used to model underspecification, speaker types represent the features of the speaker about which the addressee is uncertain, both which speech act the speaker has undertaken and their payoffs in various resulting situations—i.e., their goals. The message set consists of a range of available utterances that the speaker can select, both more and less specific. The set of addressee actions, as we will understand it, represents the set of possible ways that the addressee can update their presuppositions concerning what the speaker has asserted—they are the different ways that interlocutors can update their conception of the conversation to reflect what act they take the speaker to have undertaken. In Stalnakerian terminology, they represent different understandings of how the speaker is proposing to update the common ground.

Given this understanding of types and actions, the signalling game proceeds as follows. If I am the speaker, I send the message that maximises my expected utility, given my type and my priors about which action you will undertake once you receive it. In turn, once I send a message, you update your priors about my type according to Bayes’ rule and select an action—i.e., select how to update your presuppositions—in a way that maximizes your expected utility in light of your updated beliefs. As we saw above, signalling games of this kind—in which there are multiple meanings available for a speaker’s message, and an audience must select one under uncertainty—are sometimes known as *ambiguity* games.³⁶

To see how an ambiguity game can help to model manipulative underspecification, consider three candidate utterances, (15-a), (15-b), and (15-c).

- (15) a. I promise to keep guns out of the hands of dangerous people
(m_{bc})

³⁶For discussion of such games, see Asher and Lascarides (2013), Parikh (1991), and Parikh (2001).




Figure 1: Manipulating with Underspecification

- b. I promise to keep guns out of the hands of people who fail background checks (m_b).
- c. I promise to keep guns out of the hands of people who aren't constitutionally entitled to them (m_c).

Here, there are two utterances that express specific contents (m_b, m_c), and one that is underspecified between them (m_{bc}). Let us also assume that there are three acts that I might be undertaking. I might be promising to institute a background check policy, t_b , I might be promising to leave gun laws as lax as the second amendment allows, t_c , or I might be pried piping, t_m . For the moment we can set aside the possibility that I am underspecifying non-strategically, and mean the disjunctive content. Finally, suppose that there are three possible ways that you can update your presuppositions, i.e. three acts you can take me to have engaged in: a_b, a_c , or a_{bc} . In the first two cases you take me to have asserted something specific, while in the third case you recognize that I am deliberately underspecifying in order to manipulate you, and so treat my utterance as infelicitous.

We can see this process in action in Figure 1. Here, we begin from an initial node T at which my type— t_b, t_c , or t_m —is determined with some initial probability. I know which act I am undertaking, but you do not; you have only a prior probability distribution P over the three types. As we saw above, we can think of P as fixed by the information in the common

ground—i.e., by our shared presuppositions—at the time of the utterance.³⁷ At the first layer of nodes, labeled with an “M” (for “message”) on the left, I choose a message—either m_b , m_c , or m_{bc} . At the subsequent layer of nodes—labeled with “A” (for “action”)—you choose an action on the basis of your beliefs about me. The combination of my type, my choice of message, and your choice of action jointly determine our payoffs, which are listed in parentheses at each leaf.

Let us focus for the moment on the case in which I am pied piping, and send the underspecified message m_{bc} , represented by the rightmost main branch of the tree in Figure 1. The key feature of this case is that my uttering this message does not tell you which of the three acts I have undertaken, and so does not tell you which of the three main branches of the tree you are on. Rather, after I send m_{bc} , you will be uncertain about my type, and will have a conditional probability distribution over the three possibilities; you will calculate $P^*(t|m_{bc})$, your posterior distribution, for each t via Bayes’ rule.

In Figure 1, your uncertainty is represented by the dotted lines surrounding the nodes at which you are set to choose an action. Because you are uncertain about whether our utilities are aligned or misaligned, you are uncertain about whether I am speaking strategically or cooperatively. On the rightmost branch, our utilities are misaligned—I am speaking strategically. On the left two branches, they are not—I am being cooperative, although sending the slightly more specified message is a tiny bit costly, indicated by the payoff $1 - \epsilon$. Your uncertainty concerning which speech act I am undertaking is uncertainty about my payoffs for your various different choices. Thus, your reasoning about what speech act I have undertaken involves reasoning about my utilities, and whether I am being fully cooperative.

In the case where I am trying to manipulate you—i.e. on the rightmost branch of Figure 1—my hope is to convince you that you are in fact on one of the two main leftward branches of the tree; I aim to convince you that I have asserted something specific. I can do this by getting you to trust that I am cooperative, and would only use m_{bc} to assert either t_b or t_c . When I send m_{bc} , you compute the posterior belief $P^*(t_b|m_{bc})$ —i.e. the probability that I am promising to institute a background check, given that I utter

³⁷This is something of an idealization, for treating there as being a unique prior fixed by the context comes close to an endorsement of *Bayesian Uniqueness*. See Lederman (2015) and D’Ambrosio (manuscript) for discussion.

(15-a)—as follows:

$$P^*(t_b|m_{bc}) = \frac{P(m_{bc}|t_b) \times P(t_b)}{\sum_{t' \in T} P(m_{bc}|t') \times P(t')}$$

The presumption of cooperativity plays a crucial role in your update; it is captured in your priors over my strategies—i.e. by the credences $P(m_{bc}|t_b) \times P(t_b)$, $P(m_{bc}|t_c) \times P(t_c)$, and $P(m_{bc}|t_m) \times P(t_m)$. The more likely you think I am to send m_{bc} when t_b is actual, and the less likely you think I am to send m_{bc} when I have a manipulative intention, the more cooperative you take me to be, and the more you will increase your credence, upon receipt of m_{bc} , that I am of type t_b .³⁸

Let's suppose that you have the following priors: $P(t_b) = .45$, $P(t_c) = .45$, and $P(t_m) = .1$. This, we can suppose, captures your priors concerning the kinds of promises that speakers might make. Further, suppose that you have the following priors over my strategies: $P(m_{bc}|t_b) = .9$, $P(m_{bc}|t_c) = .2$, and $P(m_{bc}|t_m) = 1$. Together, this means that, conditional on intending t_b , your credence that I will use message m_{bc} is very high. Roughly, you anticipate that when I intend to communicate t_b , I will do so by using an underspecified message—you are trusting. However, you don't think that I will use the underspecified message to communicate t_c , and you are certain that if I am aiming to manipulate, then I will use this message. Bayes' theorem then yields the following:

$$\begin{aligned} P^*(t_b|m_{bc}) &= \frac{.9 \times .45}{(.9 \times .45) + (.1 \times .45) + (1 \times .1)} \\ P^*(t_b|m_{bc}) &= .608 \end{aligned}$$

In this situation, you will raise your credence from an initial .45 to .608 that I am promising to keep guns out of the hands of people who fail background checks. Thus, your best guess about what I mean is t_b . In turn, using this guess about what I meant, you update in a way that maximises your expected utility, where expected utility is calculated as in (16), and maximized

³⁸This component of the model is extremely similar to the RSA model. Just as here, in the RSA model, upon receipt of a message, the listener updates their beliefs about the speaker's intention in terms of their priors over the speaker's utility function. If the speaker is aiming to be fully informative, the RSA model gives the result that the listener should raise their credence in some content other than the literal content of the utterance—they should compute an implicature. The only difference here is that I am considering strategic speech, which the RSA model does not address.


as in (17):

$$(16) \quad EU_R(m, a, P^*) = \sum_{t \in T} P^*(t|m)U_R(t, m, a)$$


$$(17) \quad r(m) = \arg \max_{a \in A} EU_R(P^*, m, a),$$

where r is the function that specifies your choice of action in response to message m . The update that maximises your expected utility in this case is a_b , and so you come to presuppose that I have asserted t_b . But this is a mistake—you are deceived about what I mean, and so have updated in a way that is to my advantage and your disadvantage. In fact, I do not have intention t_b , I am pied piping, and so I am not asserting anything at all. As a consequence, I get payoff 2, and you get payoff 0; you have been manipulated; you end up presupposing that I have asserted t_b , even though I haven't.


In the case where you adopt interpretation a_b , you update your presuppositions with the information that I have promised to institute background checks. Thus, you take me to have made a certain promise. But as we saw above, in impoverished contexts, in which I know neither your background views nor your priors about how trustworthy I am, how you update your presuppositions will be opaque to me. In light of my ignorance of your priors and your utilities, I will not know which of the three updates you undertook. This will lead to fragmentation of the following sort:



Here, the dotted blue line represents the fact that I am now uncertain about how you have updated your presuppositions and so am uncertain about which new context we are in. There are three possible updates you might have undertaken— a_b , a_c , and a_{bc} , which lead to three possible defective contexts, or “context fragments” c_1 , c_2 , and c_3 , respectively. Since you have updated with a_b , c_1 is the actual context, but since the initial context was impoverished, I will be uncertain about which of c_1 , c_2 , or c_3 is actual. In impoverished contexts, underspecification leads to fragmentation.



However, this situation will often not last for long. Rather, your response—whether initially or further downstream in conversation—will typically reveal how you have interpreted me, which then “defragments” the context, rendering it non-defective. In the figure below, r_1 represents your response. Continuing with our example, we can take your response to be the utterance “Great, but what are you going to do about all the guns already in circulation?” Given this response, I now have a much better idea what you took me to mean—your utterance resolves my uncertainty about which action you undertook—you updated with a_b rather than a_c . Given this, all of the credence which I assigned to c_2 and c_3 is then transferred to c_1 , and actions a_c and a_{bc} , which led to contexts c_2 and c_3 respectively, are eliminated from the game tree—defragmentation occurs.



After this kind of defragmentation, I now know how you have updated, which makes the context whole again. The dotted lines here represent the way that my credences shift from being distributed over a_b , a_c , and a_{bc} to my being certain that you have updated on a_b . But in responding as you did, you have also *added* something to the common ground: that you are positively disposed to gun-control legislation. Given this new information, I am then at liberty to commit, retroactively, to an interpretation on which I promise to implement just such legislation by offering response r_2 , the utterance in (10-b). In uttering r_2 , I retroactively commit to having promised to institute background checks.

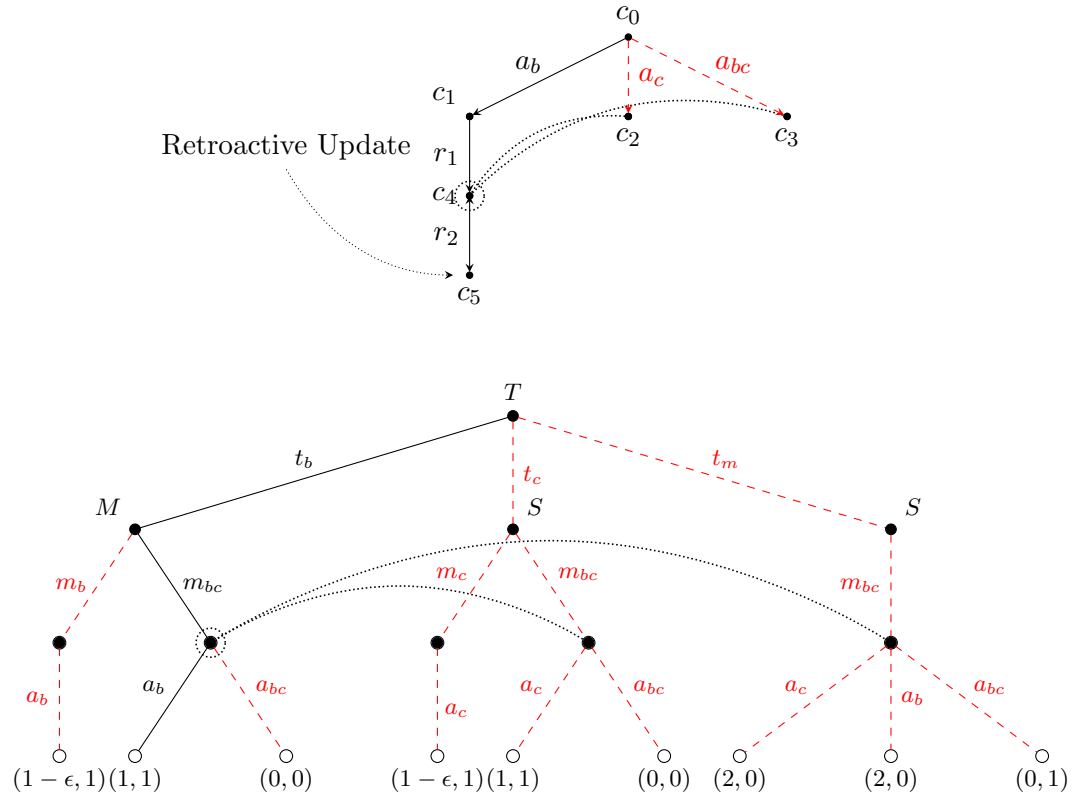


Figure 2: Manipulative Retroactive Update

In offering this response, I create a new context, c_5 , by having a retroactive effect on the conversation and the common ground. Initially you were uncertain about what I promised, and chose an interpretation given your priors. But once I offer response r_2 , your uncertainty is resolved. Conditional on this response, your credence that I undertook speech act t_a is 1. Thus, my response prunes the game tree: it eliminates branches to which you formerly assigned positive credence. Now, having undertaken act a_b , and so, with some risk, taken me to have meant t_b , you eliminate that uncertainty retroactively. My response retroactively and explicitly commits me to t_b .

Of course, had your priors been different, conversation might have gone differently. You might have, for instance, challenged my utterance, as in (11), and in so doing revealed that you disagreed with the particular interpretation you assigned. In that case, I could have retroactively committed to an alternative interpretation toward which you were more positively dis-

posed. In that case, my retroactive commitment would have eliminated different branches from your game tree: for instance, if you undertook action a_b , but rejected by assertion, my retroactive commitment would have pruned the branch on which I send message m_{bc} but have type t_b . The key feature of this dynamic is that given your uncertainty, I will always—or at least almost always—be at liberty to commit to any of the interpretations not eliminated by the context.

There is, however, an issue of just how much uncertainty needs to be involved for me to plausibly, rather than implausibly, deny the interpretation you have adopted and opt for another. Is there a clear threshold at which my denial becomes implausible? This is not something I can discuss in any detail here. But what is useful about the model just developed is that it allows us to address questions of plausible and implausible deniability as questions about how probable certain interpretations are in light of the information contained in the common ground.

Finally, I wish to return to the issue of specificity. Above we set aside the possibility that in speaking as you did, you simply meant the disjunction of the candidates left open by the context, but were not speaking strategically. But we know that fully cooperative speakers often specify only as much as their goals require, and may be indifferent beyond that (Buchanan, 2010; King, 2018, 2021; Abreu Zavaleta, 2021a). As we saw above, to account for fully cooperative underspecification, we need to expand the type and action spaces to include disjunctive contents. We can then treat the audience as having credences over both the specific contents and the disjunctive ones, along with the possibility that the speaker is speaking strategically. This generalization is relatively straightforward—the disjunctive content introduces a further type over which you have credences and a further action corresponding to your treating the disjunction as the content of my assertion.

8. Conclusion

I wish to conclude by considering some of the consequences of the picture of strategic underspecification just presented for our understanding of ordinary, non-strategic communication, and in particular for our understanding of assertion. Above I argued that, when underspecification is used strategically, the speaker deliberately puts the audience in a position where they must take a guess concerning the content of their assertion. Often, this has the effect of fragmenting the context—updating under uncertainty temporarily creates a defective context that is only resolved, or defragmented,

as conversation proceeds and the speaker overtly commits to one or another content. In strategic conversation, this dynamic affords a number of advantages to the strategic speaker, advantages I catalogued above.

But guesswork is not limited to strategic conversation. Underspecification is pervasive in non-strategic communication as well, and given that many non-strategic contexts are impoverished, it is plausible that the dynamic characteristic of strategic underspecification—guessing, temporary fragmentation, and subsequent defragmentation—is likewise common in non-strategic conversation. This should not surprise us. Linguistic communication is riddled with uncertainty about our interlocutors and the contents of their speech acts, and interlocutors must nonetheless do their best to determine the contents of speech acts and decide how to respond on that basis.

If this general picture is correct, it has important consequences for our understanding of the relationship between assertion and the common ground. Stalnaker and many who follow his lead claim that the content of an assertion is the content with which the common ground is updated when that assertion is accepted. But the ubiquity of uncertainty in communication makes this position difficult to sustain. When a speaker asserts something, that assertion can only update the common ground if every interlocutor is certain of its content. This is captured in Stalnaker's principle of *Uniformity*, which tells us that the context at the time of an utterance must *entail* a unique content for the utterance, on pain of infelicity.

Once we recognize that guesswork is a perpetual and ineliminable feature of communication, it is implausible to think of assertions as having the essential effect of directly updating the common ground. Rather, assertions may update the common ground, under certain highly idealized circumstances, but often they will not—at least immediately—and exactly what the content of an assertion was will emerge only as conversation continues and interlocutors gain new information about the speaker's intentions. Thus, the picture of strategic underspecification I have offered offers us a novel view of how assertions interact with the common ground. It is not that assertions merely add their content to the common ground upon being accepted. Rather, interlocutors will frequently have divergent views of what has been asserted, and only through further interaction can that content come to update interlocutors' store of commonly-accepted information.

References

- Abreu Zavaleta, Martín, 2021a. “Communication and Indifference.” *Mind and Language*, 36(1): 81–107.
- Abreu Zavaleta, Martín, 2021b. “Communication and Variance.” *Topoi*, 40(1): 147–169.
- Aragonès, Enriqueta and Neeman, Zvika, 2000. “Strategic Ambiguity in Electoral Competition.” *Journal of Theoretical Politics*, 12(2): 183–204.
- Aragonès, Enriqueta and Postlewaite, Andrew, 2002. “Ambiguity in Election Games.” *Review of Economic Design*, 7: 233–255.
- Armstrong, Joshua, 2013. *Language Change in Context*. Ph.D. thesis, Rutgers University.
- Asher, Nicholas and Lascarides, Alex, 2013. “Strategic Conversation.” *Semantics and Pragmatics*, 6(2): 1–62.
- Atlas, Jay, 2005. *Logic, Meaning, and Conversation: Semantical Undeterminacy, Implicature, and their Interface*. Oxford University Press, Oxford.
- Bach, Kent, 1994. “Conversational Implicature.” *Mind and Language*, 9(2).
- Bach, Kent and Harnish, Robert M., 1979. *Linguistic Communication and Speech Acts*. The MIT Press, Cambridge, MA and London, England.
- Barker, Chris, 2002. “The Dynamics of Vagueness.” *Linguistics and Philosophy*, 25: 1–36.
- Barker, Chris, 2012. “Negotiating Taste.” *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(2-3: Contextualism and Relativism): 240–257.
- Blume, Andreas and Board, Oliver, 2014. “Intentional Vagueness.” *Erkenntnis*, 79: 855–899.
- Blutner, Reinhard, 2000. “Some Aspects of Optimality Theory in Natural Language Interpretation.” *Journal of Semantics*, 17: 189–216.
- Buchanan, Ray, 2010. “A Puzzle About Meaning and Communication.” *Noûs*, 44(2): 340–371.
- Caie, Michael, 2023. “Context Dynamics.” *Semantics and Pragmatics*, 16.
- Carston, Robyn, 1988. “Implicature, Explicature, and Truth-Theoretic Semantics.” In Asa Kasher, editor, “Pragmatics: Critical Concepts,” volume 2. Broadview Press.

- Charlow, Nate, 2021. “Metasemantic Quandaries.” In Billy Dunaway and David Plunkett, editors, “Meaning, Decision, and Norms: Themes from the work of Alan Gibbard,” 171–202. Maize Books, Ann Arbor.
- Coppock, Elizabeth, 2018. “Outlook-based semantics.” *Linguistics and Philosophy*, 41: 25–64.
- D’Ambrosio, Justin, forthcoming. “A Theory of Manipulative Speech.” *The Monist*.
- D’Ambrosio, Justin, manuscript. “Clouds and Credences: A Probabilistic Approach to Underspecification.”
- D’Ambrosio, Justin and Hedden, Brian, 2024. “Multidimensional Adjectives.” *Australasian Journal of Philosophy*, 102(2): 253–277.
- Davis, Wayne A., 2019. “Implicature.” In Oxford Handbooks Editorial Board, editor, “The Oxford Handbook of Topics in Philosophy,” Oxford Handbooks. Oxford University Press.
- Deigan, Mike, 2022. “Stupefying.” *Philosopher’s Imprint*, 22(5).
- Dinges, Alexander and Zakkou, Julia, 2023. “On Deniability.” *Mind*, 132(526): 372–401.
- Dorst, Kevin and Mandelkern, Matthew, 2023. “Good Guesses.” *Philosophy and Phenomenological Research*, 105(3): 581–618.
- Drożdżowicz, Anna, 2022. “Making it precise—Imprecision and underdetermination in linguistic communication.” *Synthese*, 200(219).
- Egré, Paul, Spector, Benjamin, Mortier, Adèle, and Verheyen, Steven, 2023. “On the Optimality of Vagueness: “around”, “between”, and the Gricean maxims.” *Linguistics and Philosophy*, 45(5): 1075–1130.
- Eisenberg, Eric M., 1984. “Ambiguity as Strategy in Organizational Communication.” *Communication Monographs*, 51: 227–242.
- Fine, Kit, 1975. “Vagueness, Truth, and Logic.” *Synthese*, 30(3/4): 265–300.
- Frank, Michael C. and Goodman, Noah D., 2012. “Predicting pragmatic reasoning in language games.” *Science*, 336(6084).
- Franke, Michael, 2009. *From Signal to Act: Game Theory in Pragmatics*. Ph.D. thesis, Institute for Logic, Language, and Computation, University of Amsterdam, Amsterdam, the Netherlands.
- Franke, Michael, 2010. “Quantity Implicatures, Exhaustive Interpretation, and Rational Conversation.”

- Franke, Michael, 2013. “Game Theoretic Pragmatics.” *Philosophy Compass*, 8(3): 269–284.
- Franke, Michael and Jäger, Gerhard, 2016. “Probabilistic pragmatics: or why Bayes’ rule is probably important for pragmatics.” *Zeitschrift für Sprachwissenschaft*, 35(1): 3–44.
- Goodman, Noah D. and Frank, Michael C., 2016. “Pragmatic Language Interpretation as Probabilistic Inference.” *Trends in Cognitive Sciences*, 20(11).
- Grice, Paul, 1989. “Logic and Conversation.” In “Studies in the Way of Words,” 22–40. Harvard University Press.
- Holguin, Ben, 2022. “Thinking, Guessing, Believing.” *Philosopher’s Imprint*, 22(1): 1–34.
- Horn, Laurence, 1984. “Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature.” In Deborah Schiffrin, editor, “Meaning, form, and use in context: linguistic applications,” Georgetown University Round Table on Languages and Linguistics, 11–42. Georgetown University Press, Washington, DC.
- Horn, Laurence, 2006. *Where semantics meets pragmatics*, volume 16 of *Current Research in the Semantics/Pragmatics Interface*, chapter The border wars: A neo-Gricean perspective. Brill.
- Horwich, Paul, 1998. *Meaning*. Clarendon Press, Oxford, Oxford.
- Jaegher, Kris De, 2003. “A Game-Theoretic Rationale for Vagueness.” *Linguistics and Philosophy*, 26(5): 637–659.
- Kamp, Hans, 1975. “Two Theories About Adjectives.” In E. L. Keenan, editor, “Formal Semantics of Natural Language,” Cambridge University Press.
- Kennedy, Christopher, 2007. “Vagueness and Grammar.” *Linguistics and Philosophy*, 30(1): 1–45.
- Kennedy, Christopher and McNally, Louise, 2005. “Scale Structure, Degree Modification, and the Semantics of Gradable Predicates.” *Language*, 81(2): 345–381.
- Khoo, Justin, 2017. “Code Words in Political Discourse.” *Philosophical Topics*, 45(2): 33–64.
- King, Jeffrey C., 2018. “Strong Contextual Felicity and Felicitous Underspecification.” *Philosophy and Phenomenological Research*, XCVII(3): 631–657.
- King, Jeffrey C., 2021. *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press, Oxford and New York.

- Kirk-Giannini, Cameron Domenico, 2018. “Uniformity Motivated.” *Linguistics and Philosophy*, 41: 665–684.
- Kirk-Giannini, Cameron Domenico, 2024. “Do Not Diagonalize.” In Una Stojnic and Ernie Lepore, editors, “The Oxford Handbook of Contemporary Philosophy of Language,” Oxford University Press.
- Kocurek, Alexander, Jerzak, Ethan, and Rudolph, Rachel E., 2020. “Against Conventional Wisdom.” *Philosopher’s Imprint*, 20: 1–27.
- Krifka, Manfred, 2012. “Definitional Generics.” In Claire Beyssade Alda Mari and Fabio del Prete, editors, “Genericity,” 372–389. Oxford University Press.
- Langton, Rae, 2018. “Blocking as Counterspeech.” In Matt Moss Daniel Fogal, Daniel W. Harris, editor, “New Work on Speech Acts,” chapter 6, 144–164. Oxford University Press.
- Lassiter, Daniel and Goodman, Noah D., 2017. “Adjectival vagueness in a Bayesian model of interpretation.” *Synthese*, 194: 3801–3836.
- Lederman, Harvey, 2015. “People with common priors can agree to disagree.” *Review of Symbolic Logic*, 8(1): 11–45.
- Lewis, David, 1970. “General Semantics.” *Synthese*, 22: 18–67.
- Lewis, David, 1986. *On the Plurality of Worlds*. Blackwell Publishers, Oxford.
- Lewis, David, 2002. *Convention*. Blackwell Publishers.
- MacFarlane, John, 2016. “Vagueness as Indecision.” *Aristotelian Society Supplementary Volume*, 90: 255–283.
- MacFarlane, John, 2020a. “Indeterminacy as Indecision, Lecture I: Vagueness and Communication.” *Journal of Philosophy*, 173(11/12).
- MacFarlane, John, 2020b. “Indeterminacy as Indecision, Lecture II: Seeing Through the Clouds.” *Journal of Philosophy*, 173(11/12).
- MacFarlane, John, 2020c. “Indeterminacy as Indecision, Lecture III: Indeterminacy as Indecision.” *Journal of Philosophy*, 173(11/12).
- Mena, Ricardo, 2023. “Metalinguistic effects.” *Mind and Language*, 38: 545–565.
- Morreau, Michael, 2002. “What Vague Objects are Like.” *Journal of Philosophy*, 99: 333–361.
- Ninan, Dilip, 2011. “Semantics and the objects of assertion.” *Linguistics and Philosophy*, 33: 355–380.
- O’Connor, Cailin, 2014. “The Evolution of Vagueness.” *Erkenntnis*, 79(4): 707–727.

- O'Connor, Cailin, 2015. "Ambiguity is Kinda Good Sometimes." *Philosophy of Science*, 82(1): 110–121.
- Parikh, Prashant, 1991. "Communication and Strategic Inference." *Linguistics and Philosophy*, 14: 473–514.
- Parikh, Prashant, 2001. *The Use of Language*. CSLI Publications, Stanford, CA.
- Parikh, Prashant, 2010. *Language and Equilibrium*. MIT Press, Cambridge and London.
- Peet, Andrew, 2016. "Testimony and the epistemic uncertainty of interpretation." *Philosophical Studies*, 173: 395–416.
- Peet, Andrew, 2017. "Epistemic Injustice in Utterance Interpretation." *Synthese*, 194: 3421–3443.
- Pollock, Joey, 2023. "Testimonial Knowledge and Content Preservation." *Philosophical Studies*, 180: 3073–3097.
- Rabern, Brian, 2012. "Against the identification of assertoric content with compositional value." *Synthese*, 189: 75–96.
- Recanati, François, 2002. "Unarticulated Constituents." *Linguistics and Philosophy*, 25: 299–345.
- Roberts, Craige, 1996. "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." In Jae-Hak Yoon & Andreas Kathol, editor, "OSU Working Papers in Linguistics, Vol 49: Papers in Semantics," The Ohio State University Department of Linguistics.
- Sassoon, Galit W., 2013a. "A Typology of Multi-Dimensional Adjectives." *Journal of Semantics*, 30(3): 335–380.
- Sassoon, Galit W., 2013b. *Vagueness, Gradability, and Typicality: The Interpretation of Adjectives and Nouns*. Current Research in the Semantics/Pragmatics Interface. Brill, Leiden and Boston.
- Saul, Jennifer, 2018. "Dogwhistles, Political Manipulation, and the Philosophy of Language." In Daniel Fogal, Daniel W. Harris, and Matt Moss, editors, "New Work on Speech Acts," Oxford University Press.
- Saul, Jennifer M., 2024. *Dogwhistles and Figleaves: How Manipulative Language Spreads Racism and Falsehood*. Oxford University Press, Oxford.
- Sorensen, Roy, 1988. *Blindspots*. Oxford University Press, New York.
- Soria-Ruiz, Andrès, 101. "Evaluative and Metalinguistic Dispute." *Australasian Journal of Philosophy*, 165–181.

- Sperber, Dan and Wilson, Deirdre, 1986. *Relevance: Communication and Cognition*. Harvard University Press, Cambridge, MA.
- Stalnaker, Robert, 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press, Oxford.
- Stalnaker, Robert, 2014. *Context*. Oxford University Press, Oxford and New York.
- Stalnaker, Robert C., 1974/2002. “Assertion.” In Paul Portner and Barbara H. Partee, editors, “Formal Semantics: The Essential Readings,” 147–162. Blackwell Publishers.
- Stalnaker, Robert C., 1998. “On the Representation of Context.” *Journal of Logic, Language, and Information*, 7: 3–19.
- Stalnaker, Robert C., 2002. “Common Ground.” *Linguistics and Philosophy*, 25: 701–721.
- Stanley, Jason, 2002. “Making it Articulated.” *Mind and Language*, 17(1 and 2): 149–168.
- Stanley, Jason, 2007. *Language in Context*. Oxford University Press, Oxford and New York.
- Stanley, Jason, 2015. *How Propaganda Works*. Princeton University Press, Princeton, NJ.
- Stanley, Jason and Szabó, Zoltán, 2000. “On Quantifier Domain Restriction.” *Mind and Language*, 15(2 and 3): 219–261.
- Stokke, Andreas, 2018. *Lying and Insincerity*. Oxford University Press, London and New York.
- Szabó, Zoltán Gendler, 2020. “The Goal of Conversation.” *Aristotelian Society Supplementary Volume*, 94(1): 57–86.
- Tomz, Michael and van Houweling, Robert P., 2009. “The Electoral Implications of Candidate Ambiguity.” *American Political Science Review*, 103(1): 83–98.
- Turner, Ken and Horn, Laurence, editors, 2018. *Pragmatics, Truth and Under-specification: Toward an Atlas of Meaning*. Current Research in the Semantics/Pragmatics Interface. Brill, Leiden and Boston.
- Tye, Michael, 1990. “Vague Objects.” *Mind*, 99(396).
- van Deemter, Kees, 2010. *Not Exactly: In Praise of Vagueness*. Oxford University Press, Oxford.
- van Inwagen, Peter, 1988. “How to Reason about Vague Objects.” *Philosophical Topics*, 16: 255–284.

Weiser, Ann, 1974. "Deliberate Ambiguity." *Chicago Linguistic Society*, 10: 723–731.

Williamson, Timothy, 1994. *Vagueness*. Routledge, London and New York.

Yalcin, Seth, 2014. "Semantics and metsemantics in the context of generative grammar." In Alexis Burgess and Brett Sherman, editors, "Metasemantics: New Essays on the Foundations of Meaning," 17–54. Oxford University Press.