Communicative Uncertainty and Best-Guess Updating

Justin D'Ambrosio[1]

Draft of Jan. 10, 2024. New version in progress.

Abstract

Communication involves uncertainty, or so the empirical literature tells us. But within what is perhaps the dominant framework for theorizing about linguistic communication—the common-ground framework—rational communication cannot involve uncertainty. For an assertoric utterance to update the common ground, each interlocutor must be certain of its content, a commitment expressed by the principle of *Uniformity*. Here, I first argue that the possibility of rational probabilistic communication shows that *Uniformity* is false, as are several views of communication it motivates. I then develop a novel account of communication under uncertainty on which rational interlocutors update their presuppositions with their *best guess* about the content of a speaker's assertion, conditional on the information in the common ground. I conclude by examining which assumptions are needed to guarantee that best-guess updates are also common-ground updates, and so integrate communicative uncertainty into the common-ground framework.

1. Introduction

Like most forms of human interaction, linguistic communication involves uncertainty. In conversation, even when a speaker does her best to make her linguistic intentions manifest, her interlocutors are often not in a position to recover or comprehend exactly what it was that she intended to communicate. Rather, her interlocutors will often be uncertain about the content of her speech act. In such circumstances, interlocutors will be forced to use whatever tools and evidence are available to them to estimate, guess, or inductively infer the speaker's intended meaning. Insofar as our beliefs

[1]Arché Philosophical Research Centre, University of St. Andrews

about what our interlocutors mean are like our other empirical beliefs, this should come as no surprise.

But what should come as a surprise is that the best-known and most widely accepted framework for theorizing about linguistic communication in the philosophy of language, the common-ground framework, does not allow for communicative uncertainty. On the contrary, according to Stalnaker, rational communication by means of assertion is impossible when interlocutors are uncertain about the content of an assertion. In the common ground framework, and more specifically on the Stalnakerian account of assertion, communicative certainty is enforced by the principle of UNIFORMITY:

Uniformity An assertoric utterance must express the same proposition relative to each world in the context set. Stalnaker (2002)

In other words, for rational communication to occur, the context at the time of an assertoric utterance must *entail* that it has a particular content. It follows from this that, conditional on the information in the common ground, rational interlocutors must be certain of the utterance's content. But how can UNIFORMITY be squared with the ubiquity of communicative uncertainty? The answer is that, whenever an interlocutor is uncertain about the content of an assertion, a rule can be stated that allows interlocutors to recover an alternative assertoric content that *is* entailed by the common ground, thus eliminating communicative uncertainty. In the common-ground model of assertion, all communicative uncertainty is merely apparent.

On its face, the principle of UNIFORMITY and the certainty that it requires may seem implausible. Nonetheless, there are reasons to think that UNIFORMITY must be true. First, in the common-ground framework, communicating rationally by means of assertion consists in getting every interlocutor to take its content for granted—it consists in updating the common ground. If any interlocutor in a conversation is uncertain about the assertion's content, the thought goes, they won't know what to take for granted if they come to accept the assertion, in which case the assertion cannot update the common ground. Second, if an interlocutor accepts an assertion

and comes to take something for granted that they are not certain is the assertion's content, there is no guarantee that their update will be same as that of, or transparent to, all other interlocutors, in which case again, the assertion cannot update the common ground. If either of these thoughts are right, then UNIFORMITY, in spite of appearances, is well-motivated.

This paper has two main goals. The first goal is to show that UNIFORMITY is false, and so its implausibility is not merely apparent. My argument against UNIFORMITY begins by focusing on the credences that interlocutors have conditional on the information in the common ground. The argument then considers a situation in which a speaker utters a sentence intending to communicate that p, and, conditional on the information in the common ground, all interlocutors have a high rational credence that the speaker has asserted that p, without being certain. I argue that if interlocutors accept the assertion and update on p—i.e. they all come to take p for granted—the result is an instance of successful and fully rational communication. In short, rational probabilistic communication shows us that UNIFORMITY is false.

The second goal is to develop a probabilistic model of communication under uncertainty that explains why communication of the kind above is rational. On the view I develop, which draws on recent work by Holguin (2022) and Dorst and Mandelkern (2023), each interlocutor takes their *best guess* about what the speaker has asserted conditional on the information in the common ground, and if they accept this guess, updates their presuppositions accordingly. Unlike other views of updating under uncertainty, such as Stalnaker's own, this view makes updating sensitive to credences: how each interlocutor updates their presuppositions depends on how confident they are that a speaker has asserted a particular content. Also in contrast with other views, this proposal has the advantage of not breaking the connection between the grammatically determined content of an utterance and the content of an assertion undertaken with it.

Given this model of updating, the paper goes on to examine which assumptions are needed to guarantee that individual best-guess updates yield an update to the common ground. The first assumption is the so-called

'Common Prior Assumption' familiar from Aumann's (1976) *Agreement Theorem*. This assumption guarantees that if interlocutors update only on the information in the common ground, then their posterior distribution of credences concerning what the speaker has asserted will be identical. But a range of further assumptions are also necessary. In particular, we must assume that a number of crucial pieces of information—including the common prior itself, that the interlocutors update their credences according to Bayes rule, and that they update their presuppositions by taking their best guess—are common ground at the time of the assertion. If these things are not common ground, as is the case when the context is relatively impoverished, then interlocutors updating with their best guess will end in miscommunication—it will yield a defective context.

The rest of the paper proceeds as follows. In §2 I describe the common ground model of conversation and motivations for principle of Uniformity. In §3 I turn to communicative uncertainty and underspecification, showing how the principle of uniformity motivates various different accounts of assertoric content under uncertainty. In §4 present my argument against Uniformity, which draws on the idea that communication can be rational when context merely makes interlocutors confident that an assertion has a particular content, even if they are not certain. In §5 I introduce the good-guess framework, and in §6 I argue that under uncertainty, interlocutors should update with their best guess concerning the content of the speaker's assertion, if they accept it. In §7 I discuss the assumptions necessary to guarantee that best-guess updates yield a common ground update, and I conclude in §8 by examining the consequences of this view for our theory of assertion.

2. Rational Communication and Uniformity

In the common-ground framework, every conversation involves a set of interlocutors or conversational participants. At any point in the conversation, there is a body of information that these interlocutors share: this is the common ground. There is significant dispute over exactly how the

common ground should be defined, and we will consider competing definitions of common ground below, but at least initially we can follow Stalnaker (2014) in defining the common ground in terms of the attitude of *common acceptance*.

Common acceptance is an infinitely iterated, joint attitude that has the structure of common knowledge or common belief. Following Lederman (2015), we can specify this this structure as follows. If ϕ is an arbitrary propositional attitude, say that a set of interlocutors mutually ϕ_1s that p if and only if all interlocutors ϕ that p, and say that a set of interlocutors mutually ϕ_n that p if they mutually ϕ_1 that they mutually ϕ_{n-1} that p. We can then say that a group of interlocutors commonly ϕs that p if they mutually ϕ_n that p for every n. Thus, interlocutors commonly accept that p if they mutually accept$_n$ that p for every n.

Given this definition, we can define the common ground of a conversation c at a world w as the set of propositions commonly accepted by interlocutors in c at w. This definition shows us how the common ground is determined by the attitudes of interlocutors in a conversation. On the assumption that propositions are sets of possible worlds, we can then define the the *context set* of c at w, CS^c_w, as the intersection of every proposition in the common ground. The context set is the set of possible worlds consistent with every proposition in the common ground.

We can also use the definition of common ground just given in order to define the attitude of presupposition as follows. An interlocutor presupposes that p just in case they believe that p is common ground. Sometimes, interlocutors will be mistaken about which propositions are common ground—their beliefs about what is common ground will be false. Interlocutors may also have different presuppositions—they will have conflicting beliefs about what is common ground—in which case the context will be *defective*. Except for when I explicitly discuss the possibility that interlocutors update their presuppositions in different ways, as I do in §6, I will presume that all interlocutors are correct in their beliefs about what is common ground between them, and so that contexts, at least prior to the point at which interlocutors update on the content of an assertion, are non-defective.

The common ground is taken to play at least two important roles in conversation. The first role is metasemantic: the common ground plays the role of context with respect to which the meanings of various context-sensitive expressions are determined. Consider an example. Suppose I utter (1):

(1) This is such a drag.

Here I choose to use an expression—a bare demonstrative—whose semantic value depends on context. In trying to make myself understood, I rely on the fact that certain facts are common ground between us: for instance, that we are at a pretentious dinner party from which I want to escape. But I also rely on the fact that you accept that I accept this, since your attitude toward mine is what allows you to determine what I intend to designate with my use of 'this'. The iterative structure of the common ground explains how we successfully communicate using context-sensitive language by relying on facts that are public between us.

The second role pertains to inquiry: the common ground serves as a model of interlocutors' joint information state at any point in a conversation. As we saw above, this information state is represented by a set of possibilities—the context set—which interlocutors treat as live for the purposes of the conversation. Inquiry is typically treated as a joint activity in which interlocutors aim to determine what the world is like by narrowing down the set of possibilities that they treat as live. Insofar as speech acts, if they are accepted, change this set, the common ground can serves as a model of how conversation fits into inquiry; it models how the joint information-state of interlocutors evolves in response to their speech acts. Further, given that each basic kind of speech act updates the common ground in a distinctive way, speech acts can be characterized in terms of their effects on the context set.

According to Stalnaker, an assertion, if accepted, changes the context by adding its content to the common ground. Following Kirk-Giannini (2018), call this the *assertion rule*. According to the Assertion Rule, if I undertake an assertion whose propositional content is p, and no one objects, my

assertion updates the context by intersecting its content with the context set, and so ruling out possibilities: the updated context is $CS^c_w \cap p$. But importantly, the assertion doesn't do this directly; it updates the context by getting each interlocutor to apply the assertion rule and so update their presuppositions. For Stalnaker, rational communication by means of assertion consists in each interlocutor coming to accept—and accept that every other accepts, and accept that every other accepts that every other accepts, *etc.*—the content of the assertion, thereby updating what interlocutors commonly accept and which possibilities they treat as live.

In his classic article on assertion, Stalnaker (2002) states three conditions that must be satisfied for this form of rational communication to occur. The first is that the content of any assertion must be true in some but not all of the worlds in the context set. This captures the idea that interlocutors should neither presuppose that the sentence is true nor that its negation is true. The second is that assertions must have a truth-value at each world in the context set. If they don't, there will be no determinate way to update the common ground, and so no determinate set of possibilities that an assertion rules out.

But the condition that is most important for our purposes is the principle of Uniformity. In short, Uniformity tells us that in order for an assertion to update the common ground, and so for rational communication to occur, an assertoric utterance must express the same proposition relative to all worlds in the context set. Stalnaker in fact endorses a slightly restricted version of this principle, which says that in order for an assertoric utterance to update the common ground, it must *update the common ground* in the same way at every world in the context set. Formally, we can state this as follows, where the CS^c_w-content of u is the interesection of the content of u with the context set CS^c_w.

Uniformity In all instances of rational communication: for all assertoric utterances u, conversations c, contents p and worlds w, if speakers rationally communicate by means of u, and there exists $w' \in CS^c_w$ such that the CS^c_w-content of u at w' is p, then for all $w'' \in \mathrm{CS}^c_w$, the

CS^c_w-content of u at w'' is p. Kirk-Giannini (2018)

This principle tells us that if an assertoric utterance updates the common ground in a certain way at some world in the context set—i.e. it has a certain CS^c_w-content at that world—then it updates the common ground in that way at all worlds in the context set: it has that same CS^c_w-content at all worlds in the context set. In other words, if an assertoric utterance is to update the common ground, context must entail that it does so in a unique way.

What motivates Uniformity? There are several considerations in favour of the principle, and I cannot adequately address them all here.[2] But Stalnaker's own motivation for proposing Uniformity seems to be that that it follows from Assertion Rule—it follows from the very idea that assertions update the common ground. Kirk-Giannini spells out Stalnaker's argument for this claim as follows. The Assertion Rule tells us that assertions update the context by adding their content to the common ground. In order for this to be possible, there must be a particular content such that every interlocutor is in a position to update their presuppositions with that content. This condition is what Kirk-Giannini calls Fair Play:

Fair Play In all cases of rational communication: For all assertoric utterances u, conversations c, contents p and worlds w, s performs u at w only if every interlocutor in c is in a position to apply Assertion Rule to CG^c_w in response to u. Kirk-Giannini (2018)

The problem is that, if Uniformity is violated, it seems that at least one interlocutor will not be in a position to apply the assertion rule, and so update their presuppositions with the content of the assertion: they will be Stuck:

[2]Kirk-Giannini (2018) offers two arguments in favour of Uniformity that do not rely on Stuck. The first draws on the idea that, if one interlocutor can felicitously ask ⌜ whether ϕ⌝, then the context cannot be uniform with respect to ϕ. —fill this in . The second argument ... I will discuss these arguments further below in my argument against Uniformity.

Stuck For all conversations c, assertoric utterances u, contents p, and worlds w, w', if w' is in CS^c_w, and if p is the CS^c_w-content of u in w', and if there is some world w'' such that w'' is in C^s_w and the C^s_w-content of u in w'' is not p, then some interlocutor in c is not in a position at w to apply Assertion Rule to CG^c_w in response to u. Kirk-Giannini (2018)

Together, Fair Play and Stuck entail Uniformity. For suppose that Fair Play and Stuck are both true, and we have a violation of Uniformity. Then by Stuck, it follows that at least one interlocutor will not be able to apply the assertion rule—i.e. they will not be able to update their presuppositions with the content of the assertion. But Fair Play guarantees that in all cases of rational communication, this is not the case, which is a contradiction. Thus, if we think that Fair Play and Stuck are both required for an assertion to update the common ground, then Uniformity is required as well.

3. Underspecification and Apparent Violations of Uniformity

The problem is that, even if Uniformity is well-motivated, apparent violations of Uniformity are everywhere. There are two forms that apparent violations of Uniformity can take. The first form arises when the common ground leaves open multiple candidate updates for an assertoric utterance, and this multiplicity leaves us unable to continue the conversation—this is when underspecification in infelicitous, and leads to what Jeffrey King (2021) calls *catastrophic failure*. These are the cases that Stalnaker himself typically considers, cases in which, for instance, audience-members cannot determine the reference of a particular indexical or demonstrative.

Consider the following situation:

> [Fire]: Smith is watching Jones through a doorway. She can see the corridor in which Jones is standing, but not the room into which Jones is looking. It is common ground between Smith and Jones that the room contains either Bill or Ben and no one else.

> Bill or Ben (whoever it is) is performing a dangerous chemical experiment. Something goes horribly awry, and Jones turns to Smith and exclaims 'He is on fire!'. Hawthorne and Magidor (2009), *cf.* Kirk-Giannini (2018)

In this situation, the addressee is simply clueless about what the speaker meant. Without some rule that allows the audience to extract a content for the utterance, they will simply be at a loss as to how to update the common ground. They will be stuck.

The second form such violations can take is when context leaves open a range of contents for a particular assertoric utterance, but this does not lead to a catastrophic failure. Rather, conversation continues smoothly. Consider an example. Suppose again that we're at that dinner party. You know that I want to leave, but you don't know whether I want to leave to go to another party, leave to get ice cream, or leave to go home. Then I utter (2):

(2) I'm ready to go.

What is the content of my assertion? In this case, it's not common ground that I've asserted that I want to leave to go to another party, that I want to leave to go get ice cream, or that I want to leave to go home. Rather, given everything in the common ground, I may well mean any of these three things, or I may be completely indifferent between them. But in this case, there is no need for me to have specified any further; I have specified enough for the purposes of the conversation. This is a case of what King calls *felicitous* underspecification.

Felicitous underspecification is a pervasive feature of natural language communication. King shows that every kind of expression whose semantic value depends on context—every supplementive, as he calls them—gives rise to instances of felicitous underspecification. In cases of felicitous underspecification, like the one above, context leaves open a range of candidate propositional contents for an assertion—call this range R—and it leaves audience members uncertain about which content in R is the content of the speaker's assertion. Felicitous underspecification appears to involve a vio-

lation of UNIFORMITY, and since such underspecification is ubiquitous, so are apparent violations of uniformity. But it also appears that conversation carries on smoothly—interlocutors do not seem to be stuck. Instead, they seem to find a way to update the common ground so that conversation can proceed smoothly.

But how is the common ground updated? The first and perhaps most widely known proposal comes from Stalnaker himself. On his view, when the common ground leaves open more than one candidate content for the content of an assertion, the common ground is updated with the *diagonal proposition*. The easiest way to think of the diagonal proposition is as follows. Consider again our case from above, in which the range R of candidate contents for my assertion includes three propositions: that I am ready to go to another party, that I am ready to go get ice cream, and that I am ready to go home. The diagonal proposition is a proposition constructed from this range. Roughly, it is the proposition that, if you meant that you were ready to go get ice cream, then you're ready to go get ice cream; if you meant that you were ready to go to another party, then you're ready to go another party; if you meant that you were ready to go home, then you're ready to go home.

More formally, the diagonal proposition is constructed from an utterance's *propositional concept*. A propositional concept is a representation of the two different ways in which an utterance's truth can depend on a world. On the one hand, since an utterance's content may vary with which world it is uttered in, different worlds may yield different contents for the utterance. But on the other hand, the various contents so-fixed may also be true or false relative to those worlds. By way of illustration, consider the matrix representing the propositional concept for my utterance of 'I'm ready to go'.

$\mathfrak{A}$	a	b	c
a	T	F	F
b	F	F	T
c	F	T	F

Here, each of the rows represents the content-determining role played

by each of the worlds. Each of the columns, by contrast, represents the truth-determining role played by those worlds, given the variety of different contents the speech act might have. In the matrix, there are three worlds: a, b, and c. In world a, the content of my utterance is that I am ready to go get ice cream. In world b, the content of my utterance is that I am ready to go to another party. In world c, I mean that I am ready to go home. Now, here are the facts in those worlds: in world a, I am ready to go get ice cream—thus the content of my utterance at a is true at a. In b, I am ready to go home—thus, the content of my utterance at b is false at b. Finally, in c, I am ready to go to another party, and so again the content of my utterance at c is false at c.

Stalnaker's proposal is that, in situations where the context leaves open multiple candidate propositional contents for an assertoric utterance, the content of the speaker's assertion is the diagonal proposition formed from that utterance's propositional concept. In other words, if, as above, we construct the two-dimensional matrix from the worlds in the context set, then the utterance's content is the proposition falling along the matrix's diagonal. Thus, in this situation, my utterance eliminates worlds b and c from the context set. While this rule yields a unique content for my utterance, and so brings it back into line with UNIFORMITY, it has the consequence that the content of my assertion is divorced from the compositionally determined content of the sentence that I utter, and perhaps divorced from my intentions in uttering it.

Recently, Jeffrey King has argued against this Stalnakerian proposal and developed an alternative. King's proposal relies on a development of the common-ground framework in which context is represented by the common ground together with a *question under discussion* (QUD), which represents the current topic or subject matter of the conversation.[3] The QUD partitions the common ground into cells representing the complete answers to the

[3]The QUD was introduced by Roberts (1996), building on the work of Groenendijk and Stokhof (1990), as the centrepiece of a formal pragmatic framework oriented toward understanding discourse structure.

question, and the goal of conversation is to add information to the common ground until it entails a complete answer to the QUD. Assertions help interlocutors pursue this goal. In adding their content to the common ground, assertions eliminate candidate answers to the QUD, and bring us closer to the goal of conversation.

King claims that in cases of felicitous underspecification, where the range of candidate updates to the common ground is R, there are in fact just three candidate updates: the disjunction of the propositions in R, the conjunction of the propositions in R, and the so-called 'doesn't matter update', in which the disjunction, conjunction, and any of the candidate propositional contents update the CS in the same way. Goes on to claim that, of these three updates, interlocutors update with the weakest proposition such that no stronger proposition provides a better answer to the QUD. Thus, for King, in cases of felicitous underspecification, the content of the speaker's assertion is typically the disjunction of candidate propositions, unless it happens to be the case that their conjunction rules out more candidate answers to the QUD.[4]

While these proposals differ in the details, they have same form. Both proposals are implementation of the idea that, when context leaves open a range of candidate updates to the common ground, interlocutors engage in a piece of reasoning—something like a form of pragmatic enrichment—that allows them to extract a unique, alternative content for the assertoric utterance. On Stalnaker's view, the content is the diagonal proposition, while on King's view, the proposition is either the disjunction or the conjunction of the propositions left open by the context. This alternative update is then posited as the content of the speaker's assertion, which brings the utterance back into compliance with Uniformity and allows it to update the common ground.

[4]Stokke (2018), in his account of the lying/misleading distinction, develops a similar view on which, in cases of underspecification, what is said by an utterance is, roughly, the weakest answer to the QUD that can be constructed by disjoining propositions left open by the context.

4. Against Uniformity: Rational Credences and Underspecification

In this section I argue that UNIFORMITY itself is false, as are the proposals that it motivates. The core of my argument against UNIFORMITY comes from the idea that, rather than simply ruling out or failing to rule out candidate assertoric contents for an utterance, the common ground also serves as a shared body of evidence that, in cases of underspecification, makes some candidate contents more probable than others to be the content of the speaker's assertion. Conditional on the information in the common ground, CG, if they are rational, interlocutors will have a subjective probability distribution over the range of candidate assertoric contents in R which represents their credences concerning the content of the speaker's utterance. We can denote this conditional probability distribution as $P(R|CG)$.

Considering the credences that rational interlocutors have conditional on the information common ground is novel in two ways. First, in typical presentations of the common-ground framework, the credences of interlocutors play no role whatsoever. The common ground framework models rational communication using categorical notions like acceptance, belief, and perhaps knowledge, but it says nothing about how these notions are related to interlocutors' credences. Second, and more specifically, in the CG-framework, context is not taken to be a source of evidence that provides interlocutors with credences, which in turn guide how each interlocutor updates their information-state in response to an utterance. Given UNIFORMITY, updating is insensitive to credences, because updating rationally under uncertainty is impossible.

But any adequate account of rational communication must countenance uncertainty, and more specifically, must allow that interlocutors' conversational actions are sensitive to their credences, both about the contents of speech acts and about various background facts relevant to conversational actions. Crucially, in cases of underspecification, context will often provide interlocutors with very strong evidence that the speaker has asserted a particular content, without making it certain that they have. In such situations,

I claim, communication can be perfectly rational without constructing an alternative unique content that is uniform across the context set.

To see this, let's return to our case from above. Suppose again that I utter (2), repeated here as (3):

(3) I'm ready to go.

Recall that there are three specific contents that I might have meant that are not eliminated by the context: I might mean that I'm ready to go get ice cream (p), ready to go to another party (q), or ready to go home (r). But just because the context hasn't ruled these contents out doesn't mean that you will be indifferent between them. Suppose that it's common ground that I really like ice cream and often want to go for ice cream after dinner, but it's also common ground that I don't love parties and I rarely go to bed before midnight.

Conditional on this information, it's highly probable that what I am asserting is that I am ready to go get ice cream, although it is not certain—let's suppose that your credence $P(p|CG) = .95$, and that your credences that I mean either of the other options make up the rest of your conditional distribution. Now suppose that in virtue of this high credence, in response to my assertion, you come to presuppose that I am ready to go get ice cream—you *guess* that I mean that I want to go get ice cream, apply the assertion rule to what you guess, and update your presuppositions accordingly. Suppose further that you were right—I had a Gricean intention to communicate that I am ready to go get ice cream, and I knew the evidence would make you very confident of that, even though I knew that the context did not *entail* that I meant that.

I claim that this is an instance of fully rational communication. Based on a store of common evidence, you have a very high credence that I mean that I am ready to go get ice cream. Based on this evidence and the high credence that it justifies, you guess that this is what I have asserted, and in accepting my utterance, you come to presuppose that I am ready to go get ice cream. Moreover, this guess is correct: I did mean that I was ready to go get ice cream. Given that you have come to presuppose exactly what I in-

tended to communicate, and have done so on the basis of commonly available evidence, it is implausible to claim that we have not rationally communicated. Nonetheless, this is exactly what Uniformity entails. Given that, conditional on the information in the common ground, your credence that I meant that I was ready to go get ice cream is below 1, Uniformity entails that we did not rationally communicate. Thus, Uniformity is false.

But here the defender of Uniformity might object as follows. In the situation above, you come to *presuppose* that I am ready to go get ice cream, rather than merely accept or believe that I do. Recall that to presuppose that p is to believe that p is common ground. But, the objection goes, so long as you are uncertain what I mean, even if you are justified in accepting that p, you are not justified in believing that p is common ground between us. The response to this objection will be presented in detail in §7, but the basic idea is as follows. Given our store of commonly accepted evidence, on plausible assumptions, it will be common ground that certain guesses are better than others. If it is likewise common ground that we take our best guess, then it will likewise be common ground how we guess. Thus, I am not only justified in coming to accept that you are ready to go get ice cream, I am also rational in coming to presuppose it.

Probabilistic communication of this kind is ubiquitous. Empirical studies show us that communication nearly always involves imprecision, estimation, and uncertainty (see, e.g. Drożdżowicz (2022) and the references therein). As a consequence, it is plausible that whenever a speaker utters a sentence whose content depends on context, her interlocutors will make use of probabilistic inference to come to a determination about its content. The common ground, as a model of context, will serve as the primary source of evidence for such judgments. But it is a mistake to think that rational communication using context-sensitive language requires certainty—relatively high rational credence is perfectly sufficient. As with other areas of inquiry, determining what our interlocutors mean, when what they mean depends on knowledge of empirical facts, is going to involve uncertainty, and our model of this form of inquiry should employ our best tools for dealing with empirical uncertainty.

Given that Uniformity is false, what becomes of the proposals we saw

above concerning the content of the speaker's assertion? I claim that both of these views make incorrect predictions concerning the content of the speaker's assertion. For consider again our running example, in which I intend to communicate that I am ready to go get ice cream, and you, on the basis of the information in the common ground, correctly but probabilistically infer that I mean just that, and update your presuppositions accordingly. In that situation, it is overwhelmingly plausible that the content of my assertion is that I am ready to go to get ice cream, rather than either the diagonal proposition or the disjunction of the propositions in the range. Thus, both of those views of assertoric content make incorrect predictions: they predict that my utterance has a content that it does not have.

The argument just offered poses a dilemma for the defender of the common-ground framework. The arguments offered at the end of section 2 provied support for the conclusion that Uniformity is entailed by Assertion Rule. Given this, here is the dilemma. On the one hand, if Uniformity is necessary for updating the common ground, then common-ground updating cannot serve as a model of rational communication, because as we've seen, it appears that rational communication can proceed without it. On the other hand, if Uniformity is not necessary for updating the common ground, then the common-ground theorist lacks an account of how rational communication proceeds in its absence.

The remainder of the paper is devoted to pursuing the second horn of this dilemma by developing a new account of how rational communication proceeds under uncertainty. The foregoing arguments rely on the idea that communication can be rational when an interlocutor moves from having a high rational credence that I meant that p to presupposing that p; it relies on the idea that under uncertainty, certain *guesses* are rational, and that rational interlocutors should update their presuppositions with their best guess concerning the content of a speakers assertions. The next three sections are devoted to, (1) presenting a theory of what it is rational for a speaker to guess, (2) showing how good guesses can provide an account of how interlocutors update their presuppositions under uncertainty, and (3) arguing that best-guess updates can likewise serve as upates to the common

ground.

5. Best Guesses

In inquiry, we are often confronted with questions while being uncertain about their answers. But typically we are not totally in the dark. Rather, our evidence makes some answers more probable to be correct than others, and this is reflected in our distribution of credences over the possible answers to that question. In such situations, if someone asks us what we think the answer to the question is, we will be forced to take a guess. Which guesses are rationally permissible, given our credences? What makes a guess a good one? These questions have recently been addressed first by Holguin (2022), and then in further detail by Dorst and Mandelkern (2023), who jointly develop a theory of good guesses.

Suppose I ask you where you think Latif will go to school. He's been accepted to Yale, Harvard, Stanford, and NYU, and while you don't know anything about him in particular, you know that other students facing the same choice have decided as follows:

Yale	Harvard	Stanford	NYU
38%	30%	20%	12%

Now suppose that your credences match this evidence—how should you answer? It turns out that good guesses exhibit surprising patterns. For instance, it seems perfectly fine to guess any of the following:

(4) a. Yale
b. Yale or Harvard
c. Yale or Harvard or Stanford
d. Yale or Harvard or Stanford or NYU

On the other hand, the following guesses seem strange:

(5) a. Harvard

b. Stanford

c. NYU

d. Yale or Stanford

e. Yale or NYU.

f. Harvard or Stanford.

g. Not Yale.

h. Harvard, Stanford, or NYU.

i. Yale, and it's cold in London today.

j. Yale, or he has a birthmark on his left toe.

Here we can take the question that I ask you—i.e. where will Latif go to school?—to be the Question Under Discussion. As we saw above, in the pragmatic framework developed by Roberts (1996), the QUD is taken to partition the context set into mutually exclusive and exhaustive sets that represent the complete answers to the question. Here, answers such as 'Harvard' or 'Stanford' are complete answers to the question, whereas disjunctions of complete answers such as 'Yale or NYU' are partial answers. Guesses, then, are always partial answers to the QUD.

What separates good guesses from bad ones? Dorst and Mandelkern (2023) state five principles which they claim separate the good guesses from the bad ones, where the guesser's credences are modeled with a probability function P that is regular over the context set.

Improbable Guessing It's sometimes permissible to answer p even when $P(p) < 0.5$.

Question Sensitivity Whether p is a permissible answer depends not just on the guesser's credence in p but also in what question is being answered.

Optionality Given any question Q, for any $k : 1 \leq k \leq |Q|$, it's permissible for your guess about Q to be the union of exactly k cells of Q.

Filtering A guess about Q is permissible only if it is filtered: if it includes a complete answer p, it must include all complete answers that are more probable than p.

Fit If a guess crosscuts a complete answer, it's impermissible.

These five constraints, Dorst and Mandelkern argue, separate guesses like (4-a-d) from (5-a-j). The problem with (5-i) and (5-j) is that they violate Fit, while the rest of the bad guesses violate Filtering. Dorst and Mandelkern then argue that what separates a good guess from a bad guess is that good guesses involve optimizing a tradeoff between accuracy and informativity, and that the principles above can be derived from the right account of this tradeoff .

D&K implement this tradeoff using the idea of a guess's *answer value*, which they take from epistemic utility theory. D&M focus on what they call *Jamesian* measures of answer value, according to which an answer to a question is valuable only if it's true, and true answers to a question are more valuable the more informative they are. Informativity is measured in terms of how many answers to the QUD an answer rules out. This measure of informativity is then weighted by the credence the agent has in that answer, as follows:

Jamesian Expected Answer-Value :

$$\begin{aligned} E^J_Q(p) &= P(p) \cdot J^{Q_p} + P(p) \cdot 0 \\ &= P(p) \cdot J^{Q_p} \end{aligned}$$

Here, J is a parameter that determines the importance of informativity, while the exponent, Q_p, tells us what proportion of the answers to Q are

ruled out by p. Thus, good guesses maximise Jamesian expected answer value, but different values for J yield different weightings of informativity, and so allow more and less specific guesses to be permissible.

Complete answers, such as 'Harvard', are maximally informative, because they rule out the highest proportion of complete answers. They're obviously the best answers, if they're true. So when J is high, the best guess will be such a complete answer—high J-values prioritise informativity over accuracy, and so allow for risky guessing. By contrast, the lowest value of J is 1, which makes informativity irrelevant to how an agent guesses. In that case, the agent's best guess will be the one in which she has the highest credence. But according to D&M, there are many different permissible J-values—i.e. many different permissible weights to assign to informativity—which accounts for why Optionality is true, and so why guesses of all different strengths are permissible, as in (4).

So far we have a theory that tells us which guesses are good and which ones aren't as well as telling us why those guesses are the good oens. But why guessing important? What role does guessing play in our cognitive lives? It turns out that, according to Holduin and D&M, guessing is ubiquitous in our cognitive lives. Holguin argues that our guesses are our beliefs, and so, for Holguin, the theory of good guesses provides an account of the relationship between rational credence and belief. Dorst and Mandelkern go on claim that good guesses play two further roles. First, they claim that good guesses play a crucial role in pragmatics: on their view, we guess when we assert. In asserting, we assert our best guess at the answer to the QUD. Second, D&M claim that we guess when we reason, and the features of guessing explain why we commit the conjunction fallacy.

6. Best-Guesses and Uncertain Updating

Good guesses have yet another important application: they can explain how we communicate under uncertainty. More specifically, the theory of good guesses can be used to provide an account of how interlocutors update their presuppositions that is sensitive to credences—an account of uncertain

updating. As we've already seen, at any point in a conversation, there is some body of information that interlocutors share, represented by the context set, and also a topic of conversation, represented by a question under discussion, that partitions the context set into mutually exclusive and exhaustive cells corresponding to the question's complete answers. Within the standard QUD + Common Ground framework, assertions help us answer the QUD by adding to the common ground and so eliminating some of its candidate answers.

Before an assertion can update the common ground or eliminate any answers to the QUD, interlocutors must determine its content. But when an assertoric utterance violates Uniformity, interlocutors are uncertain about this content, and so are faced with a question—*what did the speaker assert?*—which temporarily takes precedence over the old QUD.[5] The complete answers to this new question—what did the speaker assert?—are exactly the range R of candidate updates to the common ground left open by the context. To illustrate, let's go back to our running example. I utter (6):

(6) I'm ready to go.

Given that my utterance violates Uniformity, it raises the following question:

(7) What have I asserted?

Recall that the context leaves open three possible contents for my utterance:

(8) a. I've asserted that I'm ready to go get ice cream.
b. I've asserted that I'm ready to go to another party.
c. I've asserted that I'm ready to go home.

[5]The idea that the QUD shifts from the standard QUD to the question 'what does the speaker assert?' is entailed by the very idea that speakers compute a conditional probability distribution over the utterance's candidate assertoric contents. Conditional probability distributions are computed by restricting a partition—i.e. a question—to a subset of the domain. When the candidate hypotheses concern the content of the assertion, the partition is just the question 'what is the content of the speaker's assertion?'

These three answers partition the context set into three cells, each of which represents one fully specific candidate content for my assertion. In order for my utterance to update the common ground, you need to figure out how you plan to update your attitudes if you accept what I say—you need to take a guess about what I've asserted.

As we saw above, given the information in the common ground, you will have a distribution of credences over (8-a-c), $P(R|CG)$, and according to this distribution, you are more confident that I meant (8-a) than that I meant (8-b) or (8-c). Let's say, then, that your distribution of credences over (8-a-c), conditional on the information in the common ground, is as follows:

Ice Cream	Party	Home
.8	.1	.1

Given this distribution, how should you update your presuppositions if you accept my assertion? The answer, I claim, is that you should update with your best guess about what I've asserted.

Best-Guess Updating Update with your best guess concerning the content of the speaker's assertion, if you accept it (the guess).

Given the discussion above, we can spell this out more formally as follows. Let $P(p|CG)$ be your credence that you assign to (8-a), conditional on the information in the common ground, where the common ground includes the manifest event that I've uttered 'I'm ready to go'. Given this credence, the *communicative value* of a guess i is the following:

$$P(i|CG) \cdot J^{Q_i}.$$

As we saw, Q_i is the proportion of answers that guess i rules out—it is a measure of the specificity of your guess—and $J \geq 1$ is a parameter weighting this measure of informativity. Thus, your best guess concerning the content of my assertion is the guess i that maximizes communicative value, where each guess is either a complete answer or a union of such answers.

So what is your guess? This will depend on your J-value. What do J-values represent? Here your J-value represents your estimate of the specificity of my intention. If you are confident that I meant one of the specific contents left open by the context, then your J-value will be high, and you will guess one of (8-a)-(8-c), rather than any disjunction thereof. But if you think that I am indifferent between (8-a)-(8-c), because nothing in the conversation turns on my being more specific, then your J-value will be low, and you will guess that I've asserted some disjunction of all the candidate updates left open by the context. The limit case is that you think I am completely indifferent between the contents left open by the context, and guess with $J = 1$, in which case you update on the disjuction of candidate contents.

Let's consider another example. Suppose I am a politician at a town hall, and you ask me: 'how are you going to address the problem of gun violence?' Suppose I then respond as follows:

(9) I'm going to do everything in my power to keep guns out of the hands of dangerous people.

My utterance contains at least two context-sensitive expressions for which context does not fix a unique semantic value: the quantificational expression 'everything in my power' and the multidimensional gradable adjective 'dangerous'. Focusing attention on the former, let's suppose that context leaves open the following three candidate contents:

(10)
a. I'm going to do everything in my power *consistent with the wishes of my constituency* to keep guns out of the hands of dangerous people. (p)
b. I'm going to do everything in my power *consistent with existing legal constraints* to keep guns out of the hands of dangerous people. (q)
c. I'm going to do everything in my power *including changing existing law* to keep guns out of the hands of dangerous people. (r)

In speaking as I do, and leaving open these candidate updates to the common ground, my assertoric utterance violates UNIFORMITY, and so leaves you uncertain about what I've asserted, and so uncertain about how to update the common ground.

Given this violation of UNIFORMITY, before updating the common ground and so ruling out questions to the original QUD, you need to answer the question 'what have I asserted?', which partitions the common ground into cells corresponding to each of the fully specific contents in the range R left open by the context. Here again there will be three: (10-a)-(10-c). Conditional on the common ground, you will then have a distribution of credences over these candidate contents, $P(R|CG)$, which we can suppose has the following values: $P(p|CG) = .6$, $P(q|CG) = .25$, $P(r|CG) = .15$. Thus, you are significantly more confident that in uttering (9), I mean (10-a) than either of the other two candidate meanings.

How should you update your presuppositions in response to my assertion? Again, it will depend on your J-value. If, for instance, you think that I am being deliberately vague, you may guess cautiously: you may guess with a low J-value and update with the disjunction of the three contents. But if you think I am trustworthy, and you think I intended something specific, you may guess with a high J-value, and take me to have asserted something specific. Holguin (2022) and Dorst and Mandelkern (2023) themselves are permissivists about good guesses. Given your credences, they claim, there are many different permissible ways to guess, corresponding to different J-values. If we adopt their view here, then in communication, there will be many different permissible ways to update under uncertainty.

However, it seems that while it may not be rationally impermissible to guess with different J-values, features of conversation may yet constrain the specificity of our guesses. One plausible idea, which draws on ideas introduced by King above, is that J-values are set by our conversational purposes, which are in turn understood as the purposes or goal of making an answer to the QUD common ground. In conversation, speakers specify as much as they need to given their goals, and not more. If achieving this goal requires specificity, but context leaves open a range of candidate contents,

then as long as you think I'm fully cooperative and competent, you can infer that I must have meant something specific. In such a case, you should raise your communicative J-value and guess accordingly.

To illustrate, consider our first example above. Even if the context does not fix a unique content for my utterance, you may be able to infer that that I meant something specific by reasoning about my conversational goals. Suppose, for instance, that the question under discussion is 'are you ready to go get ice cream?'. Now suppose that you utter (11):

(11) 'I'm ready!'

Given the QUD, you'll be very confident that I mean that I mean something specific. Here it clearly matters which of the more specific contents I'm asserting, and you are clearly licensed to guess that I asserted that I'm ready to go for ice cream. Importantly, however, context does not *entail* that I meant something specific, nor does it entail that I mean that I'm ready to go get ice cream. Rather, it makes both very probable. Your guess should reflect this.

But in other circumstances, it may not matter which of the various candidate updates I mean. Suppose, by contrast, that I ask you 'are you ready to leave?' and you respond with 'I'm ready'. Again, suppose that the context does not entail a unique update to the common ground. But here, given conversational purposes—i.e. answering the QUD—and given that you know I'm interested in achieving conversational purposes, you can reason to the conclusion that I must not have meant anything specific, because it wouldn't help me achieve my conversational goals. Thus, you lower your J-value, and guess that I asserted the weakest possible content: the disjunction of the three candidate contents.

Given the idea that speakers specify only as much as is conducive to their conversational purposes, we can the following principle that connects conversational J-values to the QUD:

Minimal J-value In cases of communicative uncertainty, guess with the lowest J-value such that any lower value would yield a worse answer

to the QUD.

This principle implements the idea that speakers, in general, try to assert the strongest answer to the QUD that they can. In cases where context leaves open a range of candidate contents, we can—at least in non-strategic contexts—rely on this assumption in deciding how to guess about the content of the speaker's assertion. If we know that a guess based on a low J-value—a highly nonspecific guess—would result in interlocutors updating on a content that doesn't achieve the speaker's goals, then it is an inappropriate guess.

7. Best-Guess Updating and Common-Ground Updating

Best-guess updating gives us a rule governing how interlocutors should update their own presuppositions in response to an utterance of whose content they are uncertain, in a way that is guided by the evidence in the common ground together with their J-values. But the crucial question is: when each interlocutor updates their presuppositions with their best guess concerning the content of an assertoric utterance, is the result ever an update to the common ground? Given Best-Guess Updating, can Assertion Rule be true? The answer is: yes, but only given certain assumptions about interlocutors credences, their J-values, and what is already common-ground between them.

First, recall that on the model of updating developed in the previous section, interlocutors have a conditional credence distribution over the various possible assertions that the speaker might have undertaken, and this conditional credence distribution guides how each interlocutor updates their information state. There are different ways of computing conditional probabilities, but the standard Bayesian view is that speakers update according to Bayes' Rule as follows:

$$P(i|CG) = \frac{P(CG|i) \times P(i)}{\sum_{(i' \in I)} P(CG|i') \times P(i')}$$

Here, the posterior distribution—i.e. the probability that the speaker has asserted a particular content i conditional on the information in the

common ground CG—depends on the particular audience-member's prior distribution about what the speaker is most likely to be saying: $P(I)$.

When interlocutors have different priors $P(I)$ about what a speaker has asserted, their resulting posterior distributions $P(I|CG)$ will differ. When interlocutors' posterior distributions over candidate assertoric contents differ, interlocutors may guess in different ways, although they won't always. Whether or not a difference in posterior credences will be sufficient to yield different guesses will depend on interlocutors' J-values. For example, if interlocutors J-values are set at 1, and so they give minimal weight to informativity, then interlocutors' guesses will be identical no matter how their credences are distributed, because their guess will be the disjunction of all credences in R.

When interlocutors have J-values that are greater than 1, their distribution of credences over the various candidate contents left open will matter to how they guess. In this situation, what is needed to guarantee that interlocutors end up with the same posterior distribution after conditioning on the information in the common ground is the *common prior assumption*.[6] The common prior assumption enforces the idea that different agents in a game have a unique prior probability distribution which is then conditioned on new information. For instance, the common Prior assumption is often associated with the principle of *Bayesian Uniqueness*, which says that given a particular body of evidence, there is a unique prior that interlocutors must have if they are fully rational.

One way to guarantee that the common prior assumption holds is to change something that we have been presuming. Instead of treating interlocutors as conditioning their prior distribution of credences on the common ground, which includes the manifest event of the speaker's utterance, we can treat interlocutors' priors as fixed by the information in the common ground prior to the manifest event of the utterance, and then treat this prior as conditioned on the event of the utterance. These two views are formally equivalent, but the latter gives us a principled reason to think that

[6]See Aumann (1976).

the common prior assumption holds, because it treats interlocutors' priors as uniquely fixed by a shared body of evidence.

The second assumption is that interlocutors guess with the same P-values. Guessing with the same P-values guarantees that each interlocutor places equal weight on specificity. Given that best guesses are determined by an interlocutor's credences and their P-values, sameness of posterior credence distribution and P-value guarantees sameness of guess. If, as was suggested in the last section, P-values are fixed by conversational purposes, then as long as interlocutors take the speaker to be aiming to answer the same QUD, then they will, if they are rational, guess with the same P-values.

Given that an interlocutor's best guess is determined by her credences together with her J-value, the previous two assumptions guarantee that interlocutors will guess in the same way. But even if all interlocutors make the same guess and so udate their presuppositions in the same way, this is not sufficient for their guess to likewise update the common ground. In order to guarantee interlocutors' best guess is likewise an update to the common ground, we must guarantee that their guess is commonly accepted among interlocutors, and this depends on making assumptions about what is common ground among interlocutors at the time of the assertion.

Four assumptions about what is common ground at the time of the assertion are sufficient to guarantee that the individual best-guess updates of interlocutors are also an update to the common ground:

Transparent Prior The common prior is common ground.

Transparent Bayesian Updating The fact that interlocutors update their credences according to Bayes' rule must be common ground.

Transparent Best-Guess Updating Best-Guess Updating must be common ground.

Transparent J-Values The J-value with which interlocutors guess must be common ground.

These four assumptions about what is common ground at the time of the assertion, together with an assumption of rationality and modest closure

assumptions, guarantee not only that interlocutors guess in the same way, but that the way they guess is itself common ground. Thus they are sufficient to show that each interlocutor updating their presuppositions with their best guess yields an update to the common ground rather than a defective content.

Importantly, however, these conditions are sufficient conditions—they *guarantee* that best-guess updates will yield a common ground update. But it is plausible that these assumptions can be weakened while still guaranteeing that best-guess updates are common-ground updates. The reason is that best guesses are fairly tolerant. Even if two interlocutors have different posterior credences over the candidate contents, their best guess may still be the same and may still be common ground. This is trivially the case when for both, $J = 1$. But when $J > 1$, and how interlocutors guess depends on their credences, their guesses may update the common ground even when their posterior credences diverge.

Consider the situation in which, two interlocutors have different prior distributions concerning the content of a speaker's assertion, In virtue of these divergent priors, after conditioning their priors on the information in the common ground, one comes to have credence .9 that the content of the assertion is p while the other has credence .8. Given that best guesses concern how high a credence is *relative to the other options*, it may be that these interlocutors guess in the same way. In this case, so long as their priors are common ground, and the other assumptions above hold, how they guess will likewise be common ground, and so their updates will yield an update to the common ground. I leave discussion of other weakenings of the set of assumptions above, and so investigation of exactly what is necessary for best-guess updates to yield a common-ground update, for further work.

8. Conclusion

Let us take stock. Above we saw a situation in which two interlocutors appear to communicate in a way that is fully rational when the addressee correctly, and on the basis of evidence provided by the common ground, guesses

what the speaker intends to communicate. If such communication can be rational, this presents a problem for the defender of the view that rational communication by means of assertion consists in that assertion updating the common ground. The problem is that, if Assertion Rule does entail Uniformity, then rational communication by means of assertion does not consist in updating the common ground. But if it doesn't, then we need an account of how interlocutors manage to update the common ground under uncertainty.

The view just developed aims to solve this problem by showing how best-guess updating can be used to integrate credences and Bayesian update into the common-ground framework, and also showing how, when certain assumptions are satisfied, individual best-guess updates can likewise yield an upate to the common ground. The has the intuitive consequence that whether interlocutors successfully communicate, and updating the common ground, or whether they miscommunicate, and end up in a defective context, depends on certain pieces of information being common ground at the time of the assertion itself. When these assumptions do not hold because the context is impoverished, there is no guarantee that communication will be successful.

But there is one final question to consider which has been in the background of our discussion all of this time: on the view I have proposed, what is the content of an assertion? The view of assertion underlying the above proposal is simple: interlocutors typically assert the proposition that they have the Gricean intention to communicate, while specifying only as fully as they need to for conversational purposes. For instance, in the party case above, when the QUD is 'are you ready to go get ice cream?', you may, with your utterance of 'I'm ready', assert that you are ready to go get ice cream. On the other hand, if the QUD is 'are you ready to leave'?, and context leaves open a range of things that you might be ready for, then you assert the disjunction of those propositions. In this way, my guess aims to track the specificity of your intention, and it is your intention that determine's the assertion's content.

This view of assertion has a number of advantages over the Stalnakerian

view outlined above. First, unlike the view that the content of an assertion is the diagonal proposition, my view does not divorce the content of an assertion from the grammatically determined content of the sentence the speaker utters. There is no process of "reinterpretation" occasioned by a violation of Uniformity. Rather, speakers assert what they intend to assert, and audiences take their best guess about what the speaker intends. When the common ground is rich enough, the speaker's assertion will update the common ground, but when it is impoverished, there may be miscommunication, even if slight. The result is a less idealized model of common-ground update that explains how and when defective contexts arise.

Second, the view I have developed incorporates both credences and Bayesian updating into our theory of the common ground, which until now has had no place for either. Typcially, the common-ground framework does not consider interlocutors' credences, and—partly due to the commitment to Uniformity—credences play no role in the explanation of communication. I take myself to have shown that updating under uncertainty is perfectly possible, and so shown how communication is sensitive to interlocutors' credences and what they know about each other's credences. Again, the result is a more realistic, and richer model of common-ground update under uncertainty.

References

Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics 4*(6), 1236–1239.

Dorst, K. and M. Mandelkern (2023). Good guesses. *Philosophy and Phenomenological Research 105*(3), 581–618.

Drożdżowicz, A. (2022). Making it precise—imprecision and underdetermination in linguistic communication. *Synthese 200*(219).

Groenendijk, J. and M. Stokhof (1990, August). *Partitioning Logical Space* (1.02 ed.). Leuven: ILLC/Department of Philosophy Universiteit van Amsterdam.

Hawthorne, J. and O. Magidor (2009, April). Assertion, context, and epistemic accessibility. *Mind 118*(470), 377–397.

Holguin, B. (2022). Thinking, guessing, believing. *Philosopher's Imprint 22*(1), 1–34.

King, J. C. (2021). *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford and New York: Oxford University Press.

Kirk-Giannini, C. D. (2018). Uniformity motivated. *Linguistics and Philosophy 41*, 665–684.

Lederman, H. (2015). People with common priors can agree to disagree. *Review of Symbolic Logic 8*(1), 11–45.

Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. In J.-H. Y. . A. Kathol (Ed.), *OSU Working Papers in Linguistics, Vol 49: Papers in Semantics*. The Ohio State University Department of Linguistics.

Stalnaker, R. (2014). *Context*. Oxford and New York: Oxford University Press.

Stalnaker, R. C. (1974/2002). Assertion. In P. Portner and B. H. Partee (Eds.), *Formal Semantics: The Essential Readings*, pp. 147–162. Blackwell Publishers.

Stokke, A. (2018). *Lying and Insincerity*. London and New York: Oxford University Press.