Jessica Keiser
Draft of 6/6/2019

The Limits of Acceptance

1. Introduction

Practically since its inception, the philosophy of language has been guided by a picture of language as a mechanism for cooperative information exchange. Deviant phenomena such as lying, while receiving a substantial level of interest from ethicists, have tended to be sidelined as insignificant aberrations within mainstream philosophy of language. Recent years, however, have seen an upsurge in attention to various forms of non-cooperative speech, whose theoretical importance has perhaps been made more salient by the current social and political climate. In *Lying and Insincerity*, Andreas Stokke makes a valuable contribution to this literature by providing a rigorous treatment of various phenomena falling under the umbrella of *insincere speech*, showing how they can be accommodated within a Stalnakerian framework of discourse analysis. There are many impressive features to Stokke's proposal, including his adaptation of this framework to allow sensitivity to questions under discussion, as well as its attention to distinctions between different varieties of insincere speech. A concomitant benefit of the proposal is that it brings out what is fundamentally at issue between two prominent conceptions of assertion: the Gricean and the Stalnakerian analyses. This will be my focus below.

The Gricean and Stalnakerian traditions have given rise to competing accounts of assertion, borne of a common ancestry which might be described as a project focused on reducing semantic facts to facts about the actions and mental states of language users. Stalnaker is heavily influenced by Grice, whom he cites as the inspiration behind the thesis he calls *the Autonomy of Pragmatics*: roughly, the idea that we can fruitfully theorize about the structure and function of discourse independently of the semantic properties of languages.[1] Various iterations and developments of each of these accounts may differ in detail, but the following simplified definitions will do for our purposes here:

> *Stalnakerian assertion*: an utterance made as a proposal to add the content of what is said to the common ground, understood as information that is mutually believed to be mutually accepted for the purposes of the conversation.
>
> *Gricean assertion*: an utterance made with the intention to elicit a belief in the audience toward what is said.[2]

We can roughly characterize *what is said* as the literal content of the sentence uttered by the speaker in its context (sometimes I will use "content of the utterance" as shorthand). Though many—including Stokke—offer refinements to this gloss, they won't be relevant to the following arguments.

Though neither Grice nor Stalnaker, to my knowledge, offer explicit definitions of lying, two competing theories have developed from their accounts of assertion. (For simplicity, I will refer to them as *Stalnakerian* and *Gricean,* without assuming that these individuals would endorse them.) Both understand lying in a roughly the same way: to lie is to assert what one believes to be false. But because they use different definitions of *assertion*, they deliver diverging predictions

[1] Stalnkaker (2014), p. 1.
[2] This roughly aligns with what Grice (1989) calls "utterer's occasion meaning".

on which speech acts count as lies. The key issue is that on the Stalnakerian approach, lying does not necessarily involve intent to deceive. To propose to add something to the common ground does not entail that the speaker intends to elicit a belief that content; so, when a speaker lies this need not involve an intention to deceive her audience. In contrast, the Gricean approach entails that lying involves some kind of deception. [3]

This difference is borne out in their respective treatment of bald-faced lies, understood as cases where one asserts something they believe to be false without intending to deceive.[4] Stokke argues for the superiority of the Stalnakerian account on the basis of its ability to accommodate the intuition that these are genuine lies. It is worth noting that intuitions on such fringe cases often diverge, highlighting the need for a methodology which relies more heavily on theoretical considerations;[5] nonetheless, I grant Stokke that classifying bald-faced lies as genuine lies fits best with ordinary use, and that maximal alignment with ordinary use is nice if you can get it. Thus, I adopt the assumption that classifying bald-faced lies as genuine lies is a *desideratum* of an account of lying. In what follows, however, I will question this and other predictions of the Stalnakerian account, arguing that they hinge crucially on how we sharpen our understanding of two technical terms: *assertion* and *official common ground*. I survey a number of potential precisifications, arguing that none provide a clear and non-circular metric for verifying the predictions at issue.[6] Because the options I consider are not exhaustive, it is possible for Stalnakerian theorists to provide a robust metric for testing the theory in controversial cases. My aim is to put pressure on them to do so, and to show that—until then—the Stalnakerian approach has no clear advantage over the Gricean approach.

2. Cases

To better understand what is at issue between these accounts, it will help to begin with a closer comparison of their predictions. In this section I'll discuss cases related to three broad categories of speech: bald-faced lies, non-serious speech (such as jokes and performances) and non-literal speech (such as sarcasm and metaphor):

2.1. Bald Faced Lies

As noted above, the Stalnakerian and Gricean accounts give diverging predictions about bald-faced lies. Consider this example from Carson (2006):

> *Cheating Student*: A student, who has cheated on her exam, is called into the Dean's office The student and the Dean share mutual knowledge of the student's guilt, and of the Dean's policy of failing to punish students unless they explicitly admit their guilt. When questioned by the Dean, the student says, "I didn't cheat".[7]

[3] Some versions are weaker, but most will still entail intention to deceive. For instance, the speaker may just intend the audience to believe that the speaker believes the content of the assertion—however, this still involves deception when the speaker believes this content to be false. See Grice (1989) and Bach and Harnish (1979).
[4] This is too narrow to capture ordinary use of the term; however, given that it picks out the class of bald-face lies that are at issue in this debate, it will do.
[5] Cf. Harris, this volume and Keiser (2016).
[6] My arguments elaborate on some points made in Keiser (2016). See Stokke (2017) for responses.
[7] Paraphrased from Stokke (2018) p. 18.

Though the student does not intend to deceive the Dean, she proposes that the content of her utterance be added to the stock of information mutually believed to be mutually accepted for the sake of conversation (that is, the *common ground*). Thus, on the Stalnakerian view, the utterance qualifies as an assertion, and—because its content is believed by the student to be false—a lie. This case illustrates that on the Stalnakerian view, lying does not necessarily involve an intention to deceive; because acceptance for the sake of conversation is a weaker attitude than belief, proposing to add false information to the common ground need not involve an intention to elicit a false belief.

In contrast, because the student does not perform her utterance with the intention that the Dean believe its content, the Gricean view does not classify it as an assertion. Thus, because lying is a particular type of assertion (one whose content is believed by the speaker to be false), the Gricean view predicts that the student has not lied—a result that is widely held to be counterintuitive.[8]

2.2. Non-literal speech

Griceans and Stalnakerians claim the same predictions with respect to non-literal speech; each say that such utterances are not assertions, and so do not qualify as lies when their literal content is believed by the speaker to be false. Stokke discusses a familiar example of *metaphor* from Grice (1989):[9]

> *Coffee*: A speaker utters the sentence "you are the cream in my coffee" to their lover, with the intention of communicating feelings of affection.

On the Stalnakerian approach, the speaker does not propose to add what is said to the common ground. The speech act is performed in order to communicate something else: roughly, that the speaker is fond of their lover. Therefore, the speaker fails to assert its literal content and does not lie if they believe it to be false.

The Gricean approach gives a similar prediction; the speaker does not assert the literal content, since they do not intend for their audience to believe it. Thus, they do not lie if they believe it to be false. The purpose of performing the utterance, the Gricean will agree, is to communicate that the speaker is fond of their lover. [10]

2.3. Non-serious speech

The Gricean and the Stalnakerian also claim the same predictions in cases of non-serious speech, including jokes and utterances made on stage; each claim that such utterances are not assertions and therefore not classified as lies in cases where the speaker believes what she says to be false. Stokke discusses an example of someone making a joke about the president when giving a

[8] Griceans, including Keiser (2016) and Harris (this volume) have suggested instead that the student is better understood as performing a different kind of speech act—one that is governed by institutionalized rules.
[9] See Grice (1989, p. 4)
[10] For neither the Gricean nor Stokke, however, will this metaphorical content be asserted, as it is not part of what is said by the speaker.

humorous speech at a formal banquet. I will use a real-life example from Michelle Wolf's 2018 White House Correspondent's Dinner speech:

> *Oil*: During her speech, Michelle Wolf utters the sentence "Trump is so broke he looked for foreign oil in Don Jr.'s hair."

To deal with this kind of case, Stokke introduces a modification to the Stalnakerian framework, which he suggests is independently motivated (more on this §3). Recall that common ground is defined as what is mutually believed to be mutually accepted for the sake of conversation. On the Stalnakerian picture the common ground also serves the function of storing information, keeping track of presuppositions, determining indexical content, etc. He notes that there is something like a common ground in operation during the speech which exhibits these features, but that it is in some sense unofficial. It is used temporarily for the sake of entertainment; once the speech is over, attendees will switch over to operating with a different stock of information, which we might call the *official* common ground. He uses this distinction to modify his definition of assertion: a speaker makes an assertion when she proposes to add the content of what she says to the *official* common ground. Because Wolf proposes to add the content of what she says to the *unofficial* ground, it does not qualify as an assertion. Consequently, it does not qualify as a lie, even though she (presumably) believes it to be false.

The Gricean will agree on this prediction for by now familiar reasons: Wolf does not intend for the audience to believe the content of the joke. The purpose of the utterance, rather, is to amuse (or perhaps—depending on the audience member—offend). Thus, the utterance neither qualifies as an assertion nor a lie.

In the next section I will look more closely at these predictions, but first let's take stock. Having considered three broad categories of speech, we've seen that the predictions of the Stalnakerian and the Gricean accounts agree that both non-literal and non-serious speech cannot qualify as lies, given that they fail to qualify as assertions—an intuitively compelling result, which is important for any viable theory to deliver. However, they give different predictions about bald-faced lies: According to the Gricean picture, such utterances are not genuine lies because they are not genuine assertions, while on the Stalnakerian view they qualify as both. Stokke argues for the superiority of the Stalnakerian account on the basis of its ability to deliver the more intuitive prediction in this category; thinking of these predictions in terms of a scoreboard, the Stalnakerian account of lying scores a 3/3 on the cases above, while the Gricean account trails behind with a mere 2/3.

In the remainder of the paper I take a closer look at this scoreboard. As we've seen, the results for Stokke's modified Stalnakerian account rest crucially on the technical notions of *assertion* and *official common ground*. Until we have a better understanding of these concepts, it is difficult to judge just how deep an advantage the Stalnakerian view of lying has over its Gricean competitor. Below I discuss various ways of fleshing them out, suggesting that none of them firmly establish these predictions. Since my survey is not exhaustive, it is in principle to provide definitions which justify these predictions. I hope to highlight the importance of doing so, and to show that—at least, in this stage of the game—it is less clear than Stokke suggests that the Stalknakerian theory has a substantive advantage over its Gricean opponent.

3. Assertion

In *Coffee* and *Titanic*, Stokke claims that the speakers have not lied, even though in both cases they believe the content of what is said to be false. This is uncontentiously the right result—while it is possible to mislead using non-literal speech, we do not blame speakers for lying because they believe the *literal* content of such utterances to be false. Recall that this result follows from the claim the Stalnakerian account does not predict them to be assertions. My worry is that the truth of this claim depends on how we are understanding *assertion,* which functions as technical term on this account. It is defined in terms of another technical term—common ground—which is defined in terms of the unanalyzed notion of *acceptance for the sake of conversation*. Here are the definitions, again, for reference:

> *Stalnakerian assertion*: an utterance made as a proposal to add the content of what is said to the **common ground**.
>
> *Common ground*: information that is mutually believed to be mutually **accepted for the purposes of the conversation**.
>
> *Acceptance for the sake of conversation*: unanalyzed.

In order for competitors to the Stalnakerian view to accept its predictions, these concepts must be explicated with enough substance and precision to verify those results. As will become apparent in the discussion below, understanding assertion in terms of acceptance for the sake of conversation—when the latter is taken as an unanalyzed pre-theoretical notion—is too nebulous to do much substantive theoretical work. While Stokke points to a number of characteristic features assertion which could potentially serve as a metric for verifying his results, I argue that none of them are able to do so.

3.1 Pre-theoretic Intuition about Acceptance

Suppose first that we rely on our pre-theoretic intuitions about acceptance for the sake of conversation to test the theory's predictions. While this may work in some cases, such intuitions are too unconstrained to be useful in the controversial cases at issue in the literature on lying. This comes out in Stokke's discussion of a case given by Don Fallis:

> *Deep Throat*: In a deserted parking garage in our nation's capital, a devious Deep Throat attempts to mislead a journalist by saying, "I am saying this only to you. And I am going to say it only once. If you repeat it (or say anything that presupposes it), I will deny it. The Attorney-General was behind the cover-up." (Fallis (2013), p. 350)

Fallis presents this as a counterexample to Stokke's view, claiming that Deep Throat (hereafter *DT*) lies even though he fails to make an assertion in the Stalnakerian sense. That is, DT does not propose to add what he says to what is mutually believed to be mutually accepted for the conversation, since he explicitly refuses any further discussion of it—even to the point of promising to deny any mention (or presupposition) of it. Fallis takes such behavior to be tantamount to an outright refusal to accept something for the sake of conversation. Stokke agrees with Fallis that DT is lying but denies the allegation that the case is a counterexample,

claiming that DT does indeed accept the content of his utterance for the sake of conversation. This example brings out the imprecision of the ordinary, pre-theoretical notion of acceptance for the sake of conversation, and the difficulty of using it to test predictions on cases.

Similar difficulties arise in cases where Stokke is committed to *denying* that the speaker has made an assertion. For instance, when speakers carry on with a useful metaphor, there is some pre-theoretic sense in which they are accepting its literal content for the sake of conversation. Stokke, however, denies that this behavior constitutes acceptance for the sake of conversation, thus avoiding the unwanted result that metaphorical utterances qualify as lies. I'll consider this case in more depth below; the point here is that a pre-theroetic notion of acceptance for the sake of conversation yields controversial results across a range of cases, thus failing to provide a convincing metric for verifying the predictions of the Stalnakerian account. An alternative strategy would be to point to some identifying feature of assertion which could be used to test cases; I consider several below.

3.2. Belief and inference

One feature Stokke uses to characterize assertion is the speaker's intention for the audience to believe its content and to use it in drawing inferences; he appeals to this feature to support the claim that DT does indeed make an assertion in the example above, contrary to Fallis' reading of the case. But this can't serve as a feature by which Stokke may distinguish assertions from other kinds of speech acts; it is the very feature of the Gricean characterization of assertion that he rejects, given his commitment to counting bald-faced lies as genuine lies.

Recall the bald-faced lie of *Cheating Student*. Stokke claims that in saying that she did not cheat, the student lies, entailing that she makes a genuine assertion. But her goal in performing this utterance is not to get the Dean to believe its content or to use it in drawing inferences. A crucial feature of the case is that she knows that the Dean already believes what she says to be false and does not hope or intend to change his mind. Rather, she is simply taking advantage of his policy of failing to punish students unless they explicitly admit their guilt. If intending to get the audience to believe the content of what is said and to use it in drawing inferences is an identifying feature of assertion, then the Stalnakerian account—like the Gricean account—will predict that bald-faced lies are not genuine lies.

3.3. Intention to communicate

Another feature that Stokke appeals to in characterizing assertion is its role in communication; he justifies that claim that metaphorical utterances fail to qualify as assertions by pointing out that—in contrast to cases of bald-faced lies—the speaker lacks the intention to communicate the literal content of what is said. The notion of communication, however, is no less controversial than the pre-theoretic notion of acceptance for the sake of conversation; thus, it cannot be used to provide straightforward predictions in contested cases. The term "communication" is unconstrained both in its ordinary and its theoretical use; depending on how they are understanding this term, theorists will make different predictions about whether or not the literal content is intended to be communicated in metaphorical speech. For instance, many have noted that the literal content of much non-literal speech must certainly be communicated in *some* sense,

given that it functions as a mechanism by which the non-literal content is communicated.[11] On this picture, while the literal content is not what we might call the *main point* of the utterance, it is nonetheless communicated. Thus, unless the notion of communication is made more precise, it cannot helpfully illuminate the Stalnakerian account of assertion nor provide an uncontroversial basis for confirming its predictions in cases of non-literal speech.

3.3. Support potential

The notion of support potential offers another basis for testing cases. For information to have support potential is for it to be possible for interlocutors to repeat it or to say things that presuppose it without thereby triggering surprise or requiring repair, where repair involves the audience either accommodating the content (allowing it to become common ground) or questioning/rejecting the content. Stokke appeals most strongly to support potential as a concrete way of characterizing assertion. At times, such characterizations sound definitional; for instance, he claims that "proposing to make something common ground is merely to propose that it be added to the stock of background information with support potential."[12]

However, this characterization also fails to establish the desired predictions. Fallis points out, for example, that support potential appears to be absent in the DT case. DT explicitly states that he will refuse to provide his utterance of "the Attorney-General was behind the cover-up" with support potential—i.e., if repeated or presupposed, he will force repair with utterances like:

1. "What are you talking about?"
2. "What makes you think the Attorney-General was involved?"
3. "He didn't."
4. "I never said that."[13]

Stokke claims such responses would indicate only a *pretense* that the utterance lacks support potential, referencing Stalnaker's claim that:

> If one is talking for some other purpose than to exchange information, or if one must be polite, discreet, diplomatic, kind, or entertaining as well as informative, then one may have reason to act as if the common background were different than one in fact knows it to be. (Stalnaker (1999) [1974] p. 51)

The idea is that though DT intends to add the content of what he says to the common ground (understood, on the present hypothesis, as the stock of background information with support potential), he has practical reasons to refuse to repeat/presuppose it, or to provide support to any utterances of the journalist's which do so. This, in turn, gives the journalist reasons of propriety to refrain from repeating or presupposing the information. Thus, both conversational participants have reason to act as if the common ground is different from the way they know it to be.

[11] See Davidson (1978), Camp (2006).
[12] Stalnaker (1974), p..69.
[13] Stokke (2018), p.

I have several worries for this response. First, an appeal to pretense threatens to render the account unfalsifiable. Understood behaviourally, support potential might have offered a metric for testing cases; whether something has support potential could be determined by observing the behavior of conversational participants, or considering their behavior in relevant counterfactual situations. However, the suggestion that DT's utterance has support potential in spite of the absence of its characteristic behavioural features renders this notion theoretically inert; it no longer functions as a guide to judging the predictions of the account. Perhaps one could say that though repair will in fact be requested in all relevant counterfactual circumstances, the utterance nonetheless has support potential because this repair is, in some objective sense, not *required*. However, I don't understand what this objective sense of requirement could amount to, nor how to test cases for its presence. If assertion is to be understood as proposing to add content to the stock of background information with support potential, then—if this cannot be gauged by considering the behavior of conversational participants—we need some other way to nail it down.

My second worry is that Stokke's response relies on a misreading of Stalnaker. Shortly preceding the passage, Stalnaker had introduced a first pass at characterizing presupposition in terms of common belief: "This notion of common background belief is the first approximation to the notion of pragmatic presupposition that I want to use."[14] In the passage quoted by Stoke, he had not yet refined this definition and appears to be using "common background" as shorthand for "common background belief". The refinement follows immediately after:

> . . . Let me suggest one way that the definition above needs to be qualified. . . when I talk to my barber, neither of us expects to learn anything; we are talking just to be civil, and to pass the time. If we haven't much to say, we may act as if the background of common knowledge is smaller than it really is. "Cold today, isn't it?" "Sure is, windy too." "Well, spring will be here soon.". . . When a conversation involves this kind of pretense, the speaker's presuppositions, in the sense I shall use the term, will not fit the definition I sketched above. That is why the definition is only an approximation. I shall say that one actually does make the presuppositions that one seems to make even when one is only pretending to have the beliefs that one normally has when one makes presuppositions.

Stalnaker's point is that we can accept things for the sake of conversation that are not commonly believed (by pretending that they are) and refuse to accept those that are commonly believed (by pretending that they are not). The conversation with his barber illustrates the latter case; while they have mutual knowledge of the weather conditions, they accept for the sake of conversation that they do not. Elsewhere, he uses Donellan's famous example to illustrate the former case:

> *Martini*: At a cocktail party, Alice says to Bob, "The man drinking a martini is a philosopher," knowing that the man is drinking Perrier. However, since she also knows that Bob believes that it is a martini, she believes that the best way to identify her intended referent is to use this description.

Stalnaker notes that Alice and Bob can accept for the purpose of the conversation that the philosopher is drinking the martini, even if *both* realize that it to be false; their pretense of belief

[14] Ibid., p. 49.

is rational because it is "an efficient way for them to communicate something true—information about the man who is falsely presupposed to be the man drinking a martini."[15] These cases motivate Stalnaker to weaken his account of common ground from that of common belief to that of common acceptance. The contrast he is drawing concerns what we mutually accept for the sake of the conversation vs. what we mutually believe—not, as Stokke suggests, between what we mutually accept for the sake of the conversation vs. what we pretend to accept. On Stalnaker's picture, our pretenses about what we believe *determine*—at least to some extent—what we are accepting for the sake of conversation. He makes no mention of pretense about acceptance, which—on this picture—would amount to a meta-level pretense about what we are pretending to believe. Thus, Stalnaker's remarks lend no support to Stokke's explanation, as DT neither believes that the content of his utterance is mutually believed nor pretends to believe this.

But there is a deeper worry here, which is that these passages from Stalnaker seem to undermine Stokke's results in other cases. Recall that in *Coffee*, the speaker utters to her lover

5. You are the cream in my coffee.

Stokke denies that this utterance receives support potential because it would be infelicitous for bystanders to utter things like

6. Jack doesn't realize that Mona thinks he's a dairy product.

Such an utterance would me met with surprise or require repair. However, he admits that one can felicitously carry on with a metaphor once it has been made; bystanders could felicitously utter

7. Jack doesn't realize that Mona thinks he's the cream in her coffee.

At first glance, the felicity of (7) seems to show that the content of (5) has support potential; it can be presupposed without requiring repair. However, Stokke claims that it is the *metaphorical*, rather than the literal, meaning of (5)—roughly, that the speaker is fond of her lover—that (7) presupposes. The problem with this response is that on the Stalnakerian picture, to presuppose that p is to accept that p is common ground—i.e., that it is mutually believed to be mutually accepted for the sake of conversation; if we are looking to support potential as a way to nail down what acceptance for the conversation amounts to, appealing to presupposition is of little use, given that it is defined in terms of this very notion.

Moreover, (5) and (7) appear to exemplify the kind of pretense discussed by Stalnaker; though we don't actually believe that Mona thinks Jack is the cream in her coffee, it is rational for us to talk as though we do because it allows us to communicate true information—that Mona is fond of Jack. As Camp has noted, metaphor is a particularly efficient way of communicating such information; it allows us to highlight many aspects of a subject without having to mention them individually.[16] While not every such utterance will be felicitous, this is in keeping with the explanation under consideration: (6) fails to efficiently communicate the kind of information we are interested in, given the purposes of our conversation—namely, to draw out similarities

[15] Stalnaker (2002), p. 718.
[16] See Camp (2006).

between what cream is to coffee, and to what Jack is to Mona. However, any sentence serving this purpose appears to be felicitous, e.g.

8. Mona thinks that Jack adds richness and flavor to her otherwise bitter cup
9. If Mona develops lactose intolerance, it will be bad news for Jack.

Thus, proponents of the Stalnakerian approach face two problems with respect to metaphor: First, they cannot appeal to support potential to deny that metaphorical utterances are assertions without relying on a circular characterization of that notion. Second, metaphorical utterances seem to fit neatly within the class of paradigm cases of pretense that Stalnaker uses to characterize mutual acceptance and distinguish it from mutual belief; if such pretense just *amounts* to aceptance in, e.g., *Martini,* we need some principled reason for denying that it does so in *Coffee*.

In sum, I've argued that neither the pre-theoretic notion of acceptance for the sake of conversation, nor various potential identifying features can be used to understand assertion in a way which serves as a metric for verifying the results of the Stalnakerian account of lying. In the next section I show that similar difficulties arise with respect to the notion of *official common ground*.

4. Official common ground

Non-serious speech such as joking and acting on a stage present a *prima facie* problem for the Stalnakerian picture of lying in that there seems to be some sense in which speakers are proposing to add the content of what they say to the common ground. For instance, in *Oil*, the purpose of the conversation is to entertain, and the speaker appears add the content that Trump is so broke that he looked for foreign oil in Don Jr.'s hair to the stock of information that is mutually accepted for that very purpose. Moreover, the utterance displays some of the other characteristic features of assertive content surveyed above. In particular, it's content has support potential: one could felicitously reply with utterances such as "Did he find anything?" It is also intended to be communicated—indeed, the outrage over Wolf's speech seems to confirm that she was successful in communicating the content of what she said. Stokke notes that the same can be said of an actor on the stage; the content of what she says seems to be added to some sort of common ground and displays typical features such as keeping track of anaphoric referents. However, Stokke suggests that there are two types of common ground—official and unofficial—and "we account for our intuition that the politician is not lying when she is joking by the independently motivated assumption that the common ground of fiction and jokes is seen as unofficial."[17] Only the *official* common ground matters for assertion—thus, their utterances of actors and comedians qualify neither as assertions nor lies. The problem again, I will argue, is that pre-theoretic intuition does not clearly establish that the common ground of serious speech is any more official than that of non-serious speech. I will consider other features that could potentially serve to distinguish official vs. unofficial common ground, and argue that none of them are able to confirm the desired predictions.

4.1. Pre-theoretic intuition

[17] Stokke (2018), p. 61.

There does seem to be an intuitive sense in which there can be multiple common grounds in operation at one time; the a set of information that we store relative to what is going on with respect to a play or comedic routine is distinct from the stock of information that we take to be relevant to real life. This distinction is motivated by theoretical considerations, including the fact that each of these common grounds is keeping track of a different set of presuppositions, indexical content, etc. However, it is less clear that there is theoretical motivation for dividing these different common grounds into two types—official vs. unofficial—and treating one as privileged, without appealing to belief and falling back into the Gricean picture. Stokke appears to motivate the distinction by appeal to pre-theoretic intuition, but I doubt that intuitions about which conversations are "official" can be useful in testing predictions on controversial cases. I myself would not know how to identify an official conversation, without reverting back to assumptions of common belief. Otherwise, the only identifying features that immediately spring to mind are things like formal dress and the presence of officers or other officiating bodies. But these sorts of considerations cannot serve to mark a robust and independently motivated distinction between serious and non-serious speech, for people intentionally exchange true information about the world in what many would classify as unofficial settings, and make jokes, etc. in what many would classify as very formal and official settings. Pre-theoretic intuition about what counts as official, then, does not appear to be able to provide robust support for Stokke's predictions.

4.2. Temporaility

One feature that Stokke uses to mark the distinction between official vs. unofficial common ground involves temporality; he suggests that official common ground "rules out temporarily stored information". While this characterization is more concrete, it does not appear to support the desired predictions. The underlying issue is that lying is differentiated from joking around or acting on a stage not by how *long* information is stored, but what conversationalists do with that information. There is in principle no limit for how long a joke can go on for, and information from very serious conversations may be stored only temporarily. For instance, the conversational of *Student* only store the information given by the bald-faced lie for the sake of the meeting—when the meeting is over and they return to ordinary conversational settings they will revert to talking as though the student is in fact guilty. Thus, the temporal duration of stored information does not appear to be a promising way to mark the distinction between official vs. unofficial common ground in a way that confirms the predictions of the Stalnakerian view.

4.3. Unproblematic Revocation

A final differentiating feature noted by Stokke is that information proposed to be added to unofficial common ground can be unproblematically revoked, while this is not the case with official common ground. For instance, if Wolf were accused of lying, she could unproblematically claim that she was just joking; this can be taken as an indication that her utterance was a proposal to add what was said to the *unofficial* common ground. The worry here, as Stokke notes, is that one might claim that the student could do the same thing; if accused of lying she could claim that she only said what she did in order to avoid punishment. Stokke, however, claims that the cases are not parallel because the student—unlike Michelle Wolf—could nonetheless be accused of lying. While this may be the case, the response is unsatisfactory because it threatens circularity. The Stalnakerian view purports to deliver intuitive predictions

about which utterances qualify as lies. The theory makes use of what is claimed to be an independently motivated distinction between official vs. unofficial common ground in order to deliver those predictions. This distinction cannot, in turn, be characterized by pre-theoretic intuitions about which utterances qualify as lies, or we will end up in a circle. Whether or not an utterance can be problematically revoked, then, cannot serve to characterize official common ground, if unproblematic revocation is characterized by pre-theoretic intuitions about cases.

5. Conclusion

It can be fruitful to use unanalyzed notions in a theory, up to a point. The notion of *acceptance*, in particular, has been of great use in the Stalnakerian framework of discourse analysis. However, the limits of this utility come out in controversial cases where different understandings of this notion deliver different predictions. I want to suggest that theorizing about lying, in particular, brings out the limits of how far we can fruitfully theorize with the notions of assertion—characterized in terms of *acceptance for the sake of conversation*—and *official common ground*. While they may be able to go a long way in helping us understand and systematize discourse, they cannot establish an advantage for the Stalnakerian approach to lying over the Gricean approach. The latter provides an analysis of its technical terms in a way that allows us to test its predictions; as is commonly the case, a drawback of this methodology is that such predictions are found to have counter-intuitive results in peripheral cases. While it would be desirable for a theory of lying to get the desired results with respect all three categories of speech considered above, it is arguably more important in the cases of non-literal and non-serious speech than in the case of bald-faced lies. The former are central speech act types; they are part of quotidian conversational behavior, and the folk theory of lying—while not completely constrained—seems to straightforwardly reject such cases as instances of lying.[18] In contrast, bald-faced lies are more peripheral, both in their role in ordinary speech and in their significance to the folk theory of lying. The Gricean account is able to deliver the more essential results using a systematic and testable theory.[19] In contrast, the Stalnakerian view leaves its technical terms unanalyzed, which renders the theory untestable in controversial cases. This approach is not problematic *in principle*—however it has limitations; in particular, the account cannot be convincingly shown to be superior to the Gricean account on the basis of its predictions in cases which crucially turn on a precise understanding of those very notions. I have surveyed a number of precisifications, arguing that none of them provide a clear, non-circular metric for verifying the predictions that could establish the superiority of the Stalnakerian account over that of its Gricean competitor. The survey, however, was not exhaustive, and so it is still possible to provide the sorts of definitions needed to provide a metric for testing the theory in controversial cases. This paper serves as an invitation for Stokke and other for proponents of the Stalnakerian account of lying to do so, and I look forward to the results. But—until then—I maintain that the Gricean theory is still a strong competitor.

References

Bach, Kent & Harnish, Robert M. (1979). *Linguistic Communication and Speech Acts*. Cambridge: MIT Press.

[18] I'm speaking cases where the literal content of what is said is believed by the speaker to be false. Whether someone can lie by, for instance, intentionally communicating false non-literal content is more contentious.
[19] Cf. Harris (this volume).

Camp, Elisabeth (2006). Metaphor and that certain 'je ne sais quoi'. *Philosophical Studies* 129 (1):1 - 25.

Davidson, Donald (1978). What Metaphors Mean. *Critical Inquiry* 5 (1):31-47.

Grice, H. Paul (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press.

Keiser, Jessica (2016). Bald-faced lies: how to make a move in a language game without making a move in a conversation. *Philosophical Studies* 173 (2):461-477.

Stalnaker, Robert (1974). Pragmatic Presuppositions. In *Context and Content*. Oxford University Press. pp. 47--62.

Stokke, Andreas (2017). Metaphors and Martinis: a response to Jessica Keiser. *Philosophical Studies* 174 (4):853-859.

Stokke, Andreas (2018). *Lying and Insincerity*. Oxford, UK: Oxford University Press.