

## Bald-Faced Lies: How to Make a Move in a Language Game Without Making a Move in a Conversation

### Introduction

According to the naïve, pre-theoretic conception, lying seems to be characterized by deception. Of course, lying can be unsuccessful, so deception need not *in fact* occur in order for a person to lie; rather, it seems as though it is *the intent to deceive* that is essential to lying. However, certain kinds of bald-faced lies appear to be counterexamples to this view, and many philosophers have abandoned it as a result. In this paper I will present a principled case against taking these bald-faced lies to be genuine lies, arguing that when people engage in this kind of bald-faced lying, they make a move in a language game without making a move in a conversation. Since lying is the kind of thing that one does when engaged in a conversation—in particular, when making an assertion—one is not lying unless one is so engaged when uttering something one believes to be false. This diffuses the apparent difficulty with the naïve view; bald-faced lies are not genuine instances of lying because they are not genuine instances of assertion. In the final section of the paper I will argue that there is an additional consideration in favor of the naïve view, which is that abandoning it comes at an extremely high price; alternative accounts which eschew the intent-to-deceive condition on lying are unable to distinguish lies from non-literal speech.

### 1. Bald-Faced Lies and the Intent to Deceive

A variety of acts may be classified as ‘bald-faced’ lies by ordinary language users; what they seem to have in common is that they involve a speaker saying something patently false.<sup>1</sup> A bald-faced lie may or may not involve an intention on the part of the speaker to deceive her audience. The cases that involve an intention to deceive fall neatly into the pre-theoretic

---

<sup>1</sup> Throughout the paper, I will assume that the content of what is said (or stated) is closely aligned with the literal meaning of the sentence uttered. I want to remain neutral on the details here, which I don’t anticipate will be relevant to the main arguments of the paper. The other philosophers I will be discussing have similar views, and any differences in detail will not matter for the purposes of this paper. The reader, then, should assume that ‘what is said’ (and ‘what is stated’) is being used in the same way during my presentation of others’ views.

conception of lying, however there are other cases of bald-face lying in which any attempt at deception is apparently lacking; these latter cases are taken to be counterexamples to the naïve view, which posits that lying necessarily involves an intent to deceive. To take a familiar example from the *Godfather Part 2*: “Frankie Five Angels” Pantangelli is called in as a surprise witness in a senate hearing to testify against the mob boss Michael Corleone. Frankie, who had fallen out with Michael, was under government protection and had made an agreement with government officials to tell the court of the manifold crimes for which Michael was responsible. In order to prevent this, Michael flies Frankie’s brother Vincenzo—mafioso and caretaker of Frankie’s children—over from Sicily to attend the hearings. On the day of the hearing, Frankie turns around and locks eyes with Vincenzo, whose presence reminds him that by testifying against Michael he puts the honor of his family as well as the wellbeing of his children at risk. Subsequently, Frankie surprises the government officials by going against their agreement and claiming under oath to have no knowledge of any wrongdoings committed by Michael Corleone. This is a paradigmatic example of a bald-faced lie in which there is no intent to deceive; there is mutual knowledge among the hearing attendees that Michael is guilty of murder, etc., and that Frankie has first-hand knowledge of these facts. In denying this, Frankie is not trying to deceive anyone into thinking otherwise, but rather trying to protect the honor and wellbeing of his family. Frankie’s false statement does not fit the pre-theoretic conception of lying, according to which lying is characterized by the intent to deceive; rather, it appears to undermine this *prima facie* plausible condition on lying, and many have taken this kind of case to be deeply problematic for the intuitive view. In the following discussion, I will be focusing specifically on these problematic cases of bald-faced lying—those in which there is an apparent lack of any intention to deceive. For the sake of simplicity, I will henceforward use ‘bald-faced lies’ to refer this particular problematic subset of bald-faced lies, though it should be understood that there may be other types of bald-faced lies that do not fall into this category.

There are two main strategies in the literature for dealing with the problem of bald-faced lies, each of which involves considerable difficulties. The first is to deny that bald-faced lies are in fact counterexamples to the naïve view by insisting that in these cases there really is an intention to deceive. The difficulty with this strategy is that one must explain away on a case-by-case basis the apparent lack of such an intention; moreover, the author of the alleged counterexample is perfectly within her rights to simply *stipulate* that there is no such intention—as long as there is still the

accompanying intuition that the case is an instance of lying, the proponent of the intent-to-deceive condition will have a problem to deal with. The second strategy is to take the counterexamples at face value and instead reject the intent-to-deceive condition on lying. The difficulty with this strategy is that the intent-to-deceive condition seems to be a crucial component in distinguishing lying from non-literal speech. The proponent of the second strategy, then, is faced with the challenge of finding something else that is capable of doing this work. For instance, take Grice's example of "you are the cream in my coffee", spoken by someone to a romantic partner. What is said by the speaker—namely, that the addressee is the cream in her coffee—is patently false. However, this does not intuitively strike us as a case of lying precisely because the speaker did not intend to deceive the audience with respect to that proposition. She did not intend to provide her audience with a reason to believe the literal content of the sentence uttered, or to believe that she herself believed it. Rather, she intended that the addressee, upon understanding her utterance of a sentence expressing an obviously false proposition, would come to recognize that the speaker meant something else—namely that the addressee is dear to the speaker, or something along these lines. In order to rule this out as a case of lying, we need to distinguish what was asserted from what was implicated by the speaker, and identify the speaker's communicative intentions with respect to each of these propositions. As I will argue later in the paper, recent accounts of lying which reject the intent-to-deceive condition problematically fail to mark the distinction between lies and non-literal speech.

While adopting the first strategy in defense the naïve conception of lying seems to involve insurmountable difficulties, there is another way to resist the problematic counterexamples; in the next section I will adopt a new, third strategy, arguing that when people engage in bald-faced lying, they are not making assertions. On the plausible assumption that to lie to assert something that one believes to be false, bald-faced lies are not genuine instances of lying because they are not genuine instances of assertion. After diffusing this worry about the naïve conception of lying, I will argue in the remaining section that there is an additional reason to stick with the naïve view; although many have tried to make the second strategy work, I will show that it is ultimately untenable because resulting accounts classify cases of non-literal speech (in which the literal content of the utterance is believed by the speaker to be false) as instances of lying.

## 2. Conversations vs. Games that Involve Language

### 2.1 Conversations

In what follows I will argue that in paradigmatic cases of bald-faced lying, the speaker is not performing an illocutionary act; rather, she is doing something with language other than making a move in a conversation. Before launching into that discussion, I'll say a bit about what it is to participate in a conversation, and what I take an illocutionary act to consist in. I will take on a broadly Austinian picture of speech acts, according to which an act of speech has three main components: The *locutionary act* is, roughly, to utter a sentence with its conventional meaning.<sup>23</sup> The *illocutionary act* is what one does *in* uttering that sentence, such as to make an assertion, issue a command, etc. The *perlocutionary act* is the act of doing something *by* uttering an expression, such as getting someone to close a door or to share one's beliefs. The illocutionary act is the most crucial in a conversation for the following reasons: We use language to help us navigate the world by coordinating action with other rational agents. Merely transmitting thought contents would not seem to serve this or any other significant purpose; we tend to take attitudes towards these contents, and it is by getting others to adopt attitudes towards them that we manage to influence their behavior, gain useful information, etc. So even though communication enters in at the level of the locutionary act to some extent (the transmission of a content), the locutionary act is not *in and of itself* a move in a conversation. Rather, it is in the service of the illocutionary act, which normally involves providing the addressee a reason to take certain attitude toward the transmitted content. Another lesson gleaned from Austin was that the perlocutionary act is not an act of communication, but rather the *result* of successful communication. While the speaker's aim in producing a speech act is often to produce a perlocutionary effect, successful communication takes place regardless of whether or not she succeeds. What it is to participate in a conversation, then, is to perform and understand illocutionary acts.

---

<sup>2</sup> I take it that to utter something with its conventional meaning involves uttering it with the knowledge of its conventional meaning, and with the expectation that one's audience shares this knowledge. An example of *failing* to utter a sentence with its conventional meaning would be a group of monolingual English speakers reading aloud from a book of French poetry.

<sup>3</sup> See J. L. Austin (1975). *How to Do Things with Words*. Clarendon Press.

## 2.2 Locutionary Acts and Illocutionary Acts

People use illocutionary acts to make a move in a conversation; however there are other sorts of games in which people can make a move using language. In these games, one makes a move by performing a locutionary act, though one need not—in addition—perform an *illocutionary* act. I will say more about these kinds of games in the next section, but because the locutionary/illocutionary distinction will play such a large part in my argument, I will first say a bit more (over and above the gloss given in the last section) about how I am using these terms. I take it that locutionary acts form a subset of acts of speaker meaning; they are those acts of speaker meaning that the speaker performs directly. Illocutionary acts also form a subset of acts of speaker meaning; they are those acts of speaker meaning which are performed with some kind of force—a reflexive communicative intention above and beyond the intentions that constitute speaker meaning.<sup>4</sup> I will be using the term ‘speaker-meaning’ in a slightly unorthodox way, but it should not matter for our purposes here; the reader can feel free to substitute a different term. I define speaker-meaning as follows:

*Meaning.* By uttering  $\sigma$ , U meant  $p$  iff for some audience A, U uttered  $\sigma$  R-intending that A actively attend to the thought  $p$  at least partially on the basis of her recognition of this intention.

(Where  $p$  ranges over all contents of thought, whether they be propositions or some other type of entity.)

The expression ‘R-intention’ (borrowed from Bach and Harnish) is shorthand for ‘reflexive intention’, where a reflexive intention is self-referential in the following way: a speaker has a reflexive communicative intention just in case part of her intention is that the addressee recognize the full contents of that very intention.<sup>5</sup> The resulting definition of speaker meaning is adapted from Grice’s account, but differs crucially from his (as well as from others working within a broadly Gricean framework, such as Lewis, Schiffer, and Loar) in that

---

<sup>4</sup> It follows that many instances of performatives do not qualify as illocutionary acts. Some may not be happy with this departure from Austin’s original characterization of the illocutionary act; such readers, however, can feel free to think of my use of ‘illocutionary act’ as picking out a subset of Austinian illocutionary acts.

<sup>5</sup> Bach, K. & Harnish, R. (1979). *Linguistic Communication and Speech Acts*. Mit Press.

it is less restrictive; a speaker means something merely by having the intention that her audience attend to a certain thought.<sup>6</sup> The audience need not be intended to have any additional response over and above that of having it in mind. If the speaker manages to get her audience to recognize her intentions and attend to the relevant proposition, this most basic of communicative acts—speaker meaning—has been successful. An American student reading sentences from her French language textbook aloud in order to improve her pronunciation does not *mean* anything by the sentences that she utters, since she lacks any reflexive communicative intentions in uttering them. In contrast, the actor on the stage *does* mean something by his utterance in the weak sense in which I have defined it; though he may not intend his audience to form a propositional attitude toward the content of his utterance, he does intend for his audience to respond to his utterance by coming to *attend* to this content.

Of course, given that we have an interest in sharing information and coordinating action, speakers will often have communicative intentions that go above and beyond that of speaker-meaning. For instance a speaker may utter a sentence not only with the intention that her audience *attend* to the proposition expressed by that sentence—but with the additional intention of getting the addressee to recognize the fact that the speaker believes the proposition, thereby coming to have reason to believe it herself. These additional reflexive communicative intentions constitute the illocutionary force of the utterance. They are involved in most familiar communicative settings, those which I am calling conversations.<sup>7</sup>

---

<sup>6</sup> See H. P. Grice (1989). *Studies in the Way of Words*. Harvard University Press, David Lewis (1975). Languages and language. In Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*. University of Minnesota Press, Stephen R. Schiffer (1972). *Meaning*. Oxford, Clarendon Press.

<sup>7</sup> One might worry that—given these definitions—people aren't performing illocutionary acts in contexts in which they are joking around or telling stories. This is results from the fact that I am using 'conversation' and 'illocutionary act', etc. as technical terms; I've defined them in the way that I've found to be the most theoretically useful, though this involves some departures from the way ordinary speakers and other philosophers may use them. This should not be material to the larger discussion here, because even if one wants to insist that people are performing illocutionary acts in telling jokes or stories, it is very unlikely that the candidate illocutionary act will be one that is needed for lying (such as assertion).

Not all acts of speaker-meaning are on a par. With each utterance, a speaker may be performing multiple acts of speaker meaning; however, she may use varying strategies to facilitate the hearer's recognition of the different things that she means. For a subset of the propositions meant by the speaker, the communication of these contents will be linguistically mediated in a purer sense than others. That is, in the case that a speaker means multiple things by an utterance, there will be one proposition which she intends the speaker to entertain on the basis of her utterance alone, while there will be other propositions which the hearer is intended to arrive at on the basis of other information as well—such as other propositions she recognizes to have been meant by the speaker. Taking an example from Grice, suppose a speaker—talking about a friend who has recently wronged him—utters the sentence “Jack’s a fine friend”. What is directly meant by the speaker in this case is that Jack is a fine friend. Though the speaker does not intend for the audience to believe this proposition, he does intend for the audience to *attend* to this proposition on the basis of his utterance. There is another proposition that the speaker means by this utterance, however—that Joe is *not* a fine friend. The speaker intends for the audience to attend to this second proposition on the basis of her recognition that the speaker meant the first proposition, together with clues from the context. This second proposition is indirectly meant by the speaker.

*Meaning, (Indirect Meaning):* By uttering  $\sigma$ , U meant<sub>i</sub>  $p$  iff for some audience A, there is some content  $q$  (distinct from  $p$ ) such that

1. U uttered  $\sigma$  R-intending that A attend to the thought  $q$  at least partially on the basis of her recognition of this intention.
2. U uttered  $\sigma$  R-intending that A attend to the thought  $p$  at least partially on the basis of (1).

(Where  $p$  and  $q$  range over all contents of thought, whether they be propositions or some other type of entity.)

*Direct Meaning:* Whatever is meant: if it is not meant indirectly, is it meant directly.

I will refer to all acts of direct speaker meaning as ‘locutionary acts’.

### 2.3 Language Games vs. Conversations

There are some games which involve language, but which do not thereby qualify as conversations. In these games, one can make a move in the game

using a locutionary act, without making a move in a *conversation*. A paradigmatic case of such a game would be constructing a proof. The rules of this game are different from those governing a conversation; they give restrictions on which sentences are permissible to write, given the sentences that one has already written. Making a move in the proof game involves writing a sentence down as part of a series; to win the game is to write the desired sentence as the last line in the proof, having made that—and each previous move—in accordance with the rules of the game. One does a proof using language—the act of writing down a line in a proof is a locutionary act—however it is not an *illocutionary* act. A teacher who presents a proof to her students on the blackboard does not have a reflexive communicative intention over and above the transmission of the content of the sentence she has written down. She does not assert anything, make a request, ask a question, etc.—she does not intend to provide her audience a reason to *do* anything with the content of the sentence on the basis of their recognition of her intention. (Granted, it does not seem out of the question that she could assert a line in her proof, but she need not do so in order to make a move in that game. I do think that one can use language to make a move in two games simultaneously, and will take up this issue later in the paper.)

Another paradigmatic case of a game that involves language—and one which more closely resembles a genuine conversation—is the exam game. There has been much discussion about this kind of case in the literature, grappling with the fact that answering an exam question seems like a case of speaker-meaning in spite of the speaker's lack of reflexive-intention to produce a belief in her audience. Stone and Lepore, for example, remark that

When a student answers a question on an oral exam, she surely means what she says. But how can we square this intuition with Grice's analysis? If her answer is correct, mustn't the examiner already have this information? But if her aim is for the examiner to credit her with a correct answer, isn't his response a function of the standards of the exam, not of his recognition of her intention to be judged correct?<sup>8</sup>

I think that Stone and Lepore were absolutely correct in their judgment that the students' aim is for the examiner to credit her with a correct answer—and that this intention was not a reflexive communicative intention. Granted, there was some sort of reflexive communicative intention involved here,

---

<sup>8</sup> Matthew Stone and Ernie Lepore (2010): "Against Metaphorical Meaning," *Topoi* 29(2):165-180

which was the student's locutionary act. He reflexively intended the examiner to recognize which proposition he intended her to entertain on the basis of his utterance. However, he had no reflexive communicative intention over and above that of transmitting the content which was his answer to the question. He therefore did not perform an illocutionary act having this particular content, and did not make a move in a conversation; he did, however, make a move in the *exam* game with his locutionary act. One makes a move in this game by submitting an answer to a question, the goal of the game being to submit all and only correct answers.

### 3. Bald-Faced Lies

In this section I will present a case against taking bald-faced lies to be genuine lies, arguing that—contrary to what has been assumed—bald-faced lies do not pose any particular difficulty for the view which takes the intention to deceive to be a necessary condition for lying. The argument will be laid out as follows: I will provide a definition of lying which I take to be accepted by both parties in the debate, according to which a lie must be a direct assertion. I will then provide an account of direct assertion and argue that—in spite of the misleading terminology—bald-faced lies are not genuine instances of lying because they are not instances of direct assertion. I will argue that though in many cases what the bald-faced liar is doing may look very much like assertion on the surface, it is in fact not a conversational move at all but rather a move in a different kind of language game.

In what follows I will take it for granted that to lie is to directly assert something that one believes to be false. I take this to be an uncontroversial assumption, something that members of both sides in the debate take to be true, if not true by definition.<sup>9</sup> Moreover, this assumption is justified by the fact that this definition of lying carries it a high level of *prima facie* plausibility; if one thought that a speaker could lie by performing a different speech act, it would be incumbent upon them to give an account of what

---

<sup>9</sup> Most philosophers working on this topic would say simply that to lie is to assert something that one believes to be false; my account adds the condition that the assertion must be direct. This addition is innocuous and merely reflects that I will be working with a more liberal account of assertion according to which one can assert things indirectly. Since most philosophers operate under the assumption that all assertions must be direct, adding this condition does not impose any restrictions on a definition of lying that they would not be willing to accept.

such as speech act would be. In the meantime, I will ignore alternative accounts, taking it for granted that for a speaker to lie is for her to directly assert something she believes to be false.

We can think of the dispute about whether or not lying necessarily involves the intention to deceive as running parallel to a dispute about assertion. The crux of this disagreement may be characterized as follows: on the one hand, there are those (typically Gricean) accounts which posit an intention on the part of the speaker to affect the epistemic condition of her audience in some way. A caricatured example would be an account according to which a speaker asserts  $p$  iff she utters a sentence  $s$  intending to induce that belief in the addressee. I will call such accounts—which posit an intention in the speaker to affect the epistemic state of the addressee—*epistemic accounts of assertion*. On the assumption that to lie is to directly assert something one believes to be false, the intent to deceive condition will fall out of any definition of lying which employs an epistemic account of assertion. I will be working with an account adapted from Bach and Harnish, which is more nuanced than the example given above in that the intention of the utterer is to provide her audience with a *reason* to believe something rather than to *induce* that belief in her:

*Assertion:* By uttering  $\sigma$ , U asserts  $p$  iff for some audience A

1. U utters  $\sigma$  R-intending that A attend to  $p$  at least partially on the basis of her recognition of this intention.
2. U utters  $\sigma$  R-intending that (1) will provide A with a reason to believe that U believes  $p$ .

Any assertion that is not indirect is a direct assertion, where indirect assertion is defined as follows:

*Indirect Assertion:* By uttering  $\sigma$ , U indirectly asserted  $p$  iff for some audience A, there is some content  $q$  (distinct from  $p$ ) such that

1. U uttered  $\sigma$  R-intending that A attend to  $q$  at least partially on the basis of her recognition of this intention.
2. U utters  $\sigma$  R-intending that (1) will provide A with a reason to believe that U believes  $p$ .

Given this account of direct assertion and our assumption that to lie to directly assert something that one believes to be false, it follows that lying must involve an intention to deceive. The kind of deception involved is more subtle than that of intending to *induce* a false belief in someone. Perhaps some of the resistance to the intent to deceive condition on lying might be due to an assumption that the intended deception would have to take the stronger form, which would be more naturally linked (given our account of lying as asserting something one believes to be false) to a much stronger epistemic account of assertion: one according to which the speaker asserts something intending to get the audience to believe it.<sup>10</sup> This kind of account appears to be too strong, however, and is open to counter-examples in any case where one asserts something without hope of convincing her interlocutor—whether the speaker herself believes that proposition to be true or false. But an epistemic account of assertion needn't be so strict; the intent to deceive condition falls out of the more nuanced account of assertion as well. I take it that if you intentionally give someone a reason to believe something you know to be false, this is an act of deception. You are intentionally worsening their epistemic condition, even if you don't actually cause (or even intend to cause) them to believe something false.

According to my account of assertion, if a speaker—in uttering something she believes to be false—reflexively intends her utterance to provide the audience a reason to believe that she believes that proposition, she will have asserted it and will have lied. Many speech acts which may be called 'bald-faced lies' will fit this description. The speaker may understand that the reason she provides her audience will be overridden by that person's other evidence—and that therefore her audience may not ultimately believe the proposition in question. But this kind of case is no counter-example to the intent to deceive condition on lying, since the utterer still intends to deceive her audience by intentionally worsening her epistemic state. The audience will be in a worse epistemic state on account of having a reason to believe something false, regardless of whether or not she actually comes to believe it. The kind of case which would be a counterexample to the naïve view is

---

<sup>10</sup> An account according to which the speaker asserts something intending to get the audience to believe that she (the speaker) *herself* believes it also qualifies as a stronger epistemic account of assertion, since it involves the intention to induce a belief in someone. My comments about the strong epistemic account of assertion that I have characterized above will apply *mutatis mutandis* to this kind of account as well.

one in which there is no attempt at deception whatsoever. But given our definition of lying and of direct assertion, no such counterexample could arise; there can be no lie without an attempt to deceive, because it follows from one's asserting something that one believes to be false that one intends to worsen the epistemic state of her addressee.

What, then, of those bald-faced lies in which there is absolutely no attempt at deception? I want to suggest that many of these cases are instances of a move made in a language game without being moves in a conversation. Failing to be conversational moves—and so failing to be assertions in particular—they are not genuine lies. Take again the example of Frankie Five Angels, who says something false on the witness stand without any intention to deceive his audience. Frankie performs a locutionary act as a move in a language game—in this case a courtroom game—but does not perform any identifiable illocutionary act and so does not make a move in a conversation. His locutionary act consists of transmitting a certain proposition—that he is unaware of the defendant having committed any crime. There is no reflexive communicative intention over and above the transmission of this content—he does not intend his audience to *do* anything with that content on the basis of the recognition of his intention—so he does not perform an illocutionary act or make a move in a conversation. But by performing this locution under oath and in this setting, he makes a move in a different game—the courtroom game. Though I won't begin to try to characterize the rules of this complex game, it is apparent that if a witness testifies that she observed a defendant commit a crime, the courtroom rules dictate that this results in a setback for the defendant. So in the case we are considering, in performing this locutionary act as a move in the courtroom game, the witness fails to take the opportunity to score points *against* the defendant. Indeed this reveals one of his motivations for making the move he did—he was afraid that if he scored points against the defendant in the courtroom game, the defendant would take revenge.

Presumably the courtroom game—and perhaps any kind of language game—is derivative from genuine conversation. It is for this reason that moves made in a courtroom game appear so much like real illocutionary acts on the surface, and why Frankie's statement has the feel of an assertion even in the case that the requisite intentions are lacking. A move made in a courtroom game—though not a real illocutionary act—is a mock illocutionary act. The courtroom game is modeled after real conversation. In fact, the courtroom game presumably evolved slowly from what were at one point genuine conversations. That is, initially people probably attempted to resolve

conflict and establish justice by coming together to have a genuine conversation about who was at fault and what to do about it. But as this procedure became more and more structured, it evolved into a game that—while appearing an awful lot like conversation—actually departs from it in important respects, having its own independent goals and rules.

Another reason why moves in the courtroom game can look so much like conversational moves is that it is possible for there to be both a genuine conversation and a courtroom game going on simultaneously—and for these games to pull together and come apart at various moments in time. In fact, it is very likely that there will be a conversation going on during most moments of the courtroom game, and that most of the speech acts performed in this context will constitute a move in both games. For instance, when answering innocuous questions—about his name and his age, perhaps—the witness on the stand is likely to have a reflexive communicative intention over and above that of merely transmitting a content. He is likely to genuinely intend to provide a reason for the questioner (or the courtroom audience in general) to believe that he believes the content of his utterance. However, this *need* not be so. And as I have argued above, the bald-faced lie of Frankie Five Angels is just such a case in which the two games come apart; he makes a move in the courtroom game with this particular locution without making a move in the conversation.

## **4. Problems with Alternative Accounts**

In this section I will look at several alternatives to the naïve account of lying, according to which the intention to deceive is a necessary condition. These alternative accounts eliminate this condition by employing non-epistemic accounts of assertion. I will argue that these accounts of assertion have trouble distinguishing genuine assertion from non-literal speech. Therefore, in the special case in which the content of a speaker's utterance is believed by her to be false—including the case in which she is merely speaking non-literally—it will follow from such an account that the speaker has lied.

### *4.1 Warranting Contexts*

The first alternative account of lying that I will consider is one that eliminates the intent to deceive condition by employing a non-epistemic account of assertion that appeals to warranting contexts. According to this account,

S asserts  $p$  iff she states  $p$  in warranting context,

where a warranting context is a context governed by a convention according to which any speaker in the context warrants the truth of what she states. To warrant the truth of one's statement is to provide a promise or guarantee that its content is true in uttering it.<sup>11</sup>

For simplicity, I will assume in this section that to lie is to assert something that one believes to be false. The directness condition employed in the last section will be captured by the fact that the particular accounts of assertion considered in this section require that the content of the assertion must be stated, which adequately captures the directness condition on lying; I take it, then, that this definition of lying does not differ in any important respect from the definition offered in the last section.

Given this account of assertion it follows that

S lies just in case S states  $p$  in a warranting context, believing  $p$  to be false.

However, whether or not one is in warranting context is orthogonal to whether or not one uses non-literal speech; a graduate student may use both literal and non-literal speech when answering a question during her oral examinations, and a stand-up comedian may use both literal and non-literal speech onstage during her act. It is problematic to suppose that the context would determine whether a speaker warrants the truth of what is stated in particular rather than what is communicated full stop. A context can be a warranting context in the sense that we are mutually expecting and trusting one another to be communicating truths, and yet speakers are free to use non-literal speech in order to accomplish this. I may use non-literal speech to knowingly communicate something, in which case there is a sense in which I have violated my interlocutor's trust and the norms we believed to be

---

<sup>11</sup> See Carson, Thomas L. (2006). The definition of lying. *Noûs* 40 (2):284–306 and Saul, Jennifer Mather (2012). *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press.

governing our conversation, even though it was clear that I didn't warrant the truth of the literal content of my utterance. I warranted the truth of *something*—namely the implicated content—and this is enough to constitute such a violation. Furthermore, I may use non-literal speech in a warranting context and—contrary to the account above—it won't be the case that I will have lied. Insofar as the context is relevant, it may play some part in determining whether or not I warrant the truth of what I communicate, whether that be directly or indirectly. Holding everything else fixed, however, whether or not I've lied or simply spoken non-literally seems to be determined by my communicative intentions. Therefore, the appeal to warranting contexts is not adequate to mark the distinction between lies and non-literal speech.

One may wonder if this worry could be allayed by appealing to the *intention* to warrant the truth of one's statement rather than by appealing to a warranting context. That is, consider the following revision to the warranting account of assertion:

S asserts that  $p$  iff in stating  $p$ , S intends to warrant the truth of  $p$  to her audience.

In this case it will follow that S lies just in case in stating  $p$  she intends to warrant the truth of  $p$  to her audience, while believing  $p$  to be false.

Though this account manages to evade the worries expressed above, it is far from clear whether it will allow for cases in which one can lie without having an intention to deceive. To warrant the truth of one's statement is to make a kind of promise or guarantee that its content is true. Standardly, however, when we make a promise it is a promise to ensure that a certain state of affairs will obtain. In the felicitous case, the promisor both believes herself to be in control over whether or the state of affairs will obtain, and has the intention to make it the case that it does. This element of control seems to be undergirding the convention of promising in some way; the reason promises are able to be taken seriously is that given the promisor's ability and intention to make it the case that a certain state of affairs obtains, the promisee has reason to expect that she will do so. In contrast, one does not standardly have control over whether or not the contents of statements are true. (Unless we have the power to ensure that the state of affairs which makes the content of the statement true obtains—but I am assuming that we standardly assert propositions that do not fall into this category.) How, then,

could there develop a convention of making such guarantees and of taking such guarantees from others seriously? Presumably the only reason that such a guarantee would carry any weight with an addressee is that the speaker putatively believes what she is saying. In this case, the addressee may reason as follows: *if the speaker believes what she is saying, then she must have reasons for so doing. Though I don't know what these reasons are, I can assume that they exist and that they would also be reasons for me.* So the guarantee derives its force from the fact that the speaker believes what she is saying. But how could one be in a position to guarantee the truth of something she did not believe to be true? She might be able to do this by leading the addressee into thinking that she believed the truth of her statement, but this would involve deception. If one intends to warrant a statement while believing it to be false, it follows that she intends to deceive her audience in the weak sense discussed in the last section; she will have intentionally provided her audience a reason to believe something that she believes to be false. So the intent to deceive condition on lying seems to fall out of an account of assertion that is defined in terms of the intention to warrant the truth of what is stated, where warranting a statement's truth is glossed as providing a promise or guarantee of its truth.

#### 4.2 Common ground account of assertion:

Another view which eliminates the intention to deceive condition on lying is that which employs a common ground account of assertion:

S asserts  $p$  iff in stating  $p$  S proposes to add  $p$  to the common ground,

where  $p$  is common ground in a group if all members are mutually believed to accept  $p$  for the purpose of the conversation.<sup>12</sup>

One may define common ground in various ways, but it is important to note that a stronger account which appeals to belief rather than acceptance as the relevant attitude will be too strong. If what it is to be part of the common ground is that all members of the conversation mutually believe  $p$ , then the resulting account of assertion will be an epistemic account of assertion and will entail an intent to deceive condition on lying. If a common ground account of assertion is to provide the resources for an alternative to

---

<sup>12</sup> See Stokke, Andreas (2013). Lying and Asserting. *Journal of Philosophy* 110 (1):33-60.

the naïve conception of lying, it must employ a weaker propositional attitude—such as acceptance—in its definition of common ground.

The problem with this proposal, however, is that acceptance is too weak a notion to ground the distinction between lying and non-literal speech. Since acceptance of  $p$  does not involve believing  $p$ , it is unclear how the provided definition is going to be able to rule out non-literal speech as cases of lying. People can carry on with saying things that are mutually believed to be false, for the purpose of conversation. Non-literal speech exemplifies this kind of behavior; I may use a metaphor to communicate something true, and my interlocutor may respond by playing along with that same metaphor for the course of the conversation. And as Stalnaker and others have pointed out, one or both parties to the conversation may continue to refer to a man as “the man drinking the martini”, knowing full well that he is drinking water; as long as both conversational participants understand which man is being referred to, this behavior won’t adversely affect the information exchange and it may in fact be distracting to make the correction. However, the definition above wrongly counts both of these cases—where speakers are developing a metaphor, or going along with something they know to be false for the sake of conversational ease and efficiency—as instances of lying. In order to deal with this problem, Andreas Stokke has introduced a distinction between official vs. unofficial common ground; the unofficial common ground is characterized by temporary acceptance, while the official common ground is characterized by permanent acceptance. However, this distinction seems to be insufficient to explain the difference between conversations which involve assertion and conversations in which it is mutually understood that we are carrying on a kind of pretense; my carrying on a pretense with you is not a function of the length of time that I want us to accept what is being said, but rather what kind of attitude we take to what is being said. We could carry on a pretense for the entire conversation without ever asserting the contents of our utterances. Similarly, it seems that we could very well assert things with the knowledge that they will only ever be temporarily mutually accepted by the participants of the conversation. The notion of acceptance, then, is too weak to ground the distinction between literal and non-literal speech.

## 5. Conclusion

I hope to have diffused a major worry for the naïve conception of lying—which takes the intention to deceive to be a necessary feature of this act—by providing a principled case against taking bald-faced lies to be genuine instances of lying. I have argued that while these acts paradigmatically constitute moves in a language game, they are not moves in a conversation. Since lying is the kind of thing that one does when engaged in a conversation—in particular, when making an assertion—one is not lying unless one is so engaged when uttering something one believes to be false. The proponent of the intent-to-deceive condition on lying, then, is justified in resisting the alleged counterexamples to her theory, which otherwise vastly outstrips her opponents' in its intuitive appeal and its ability to account for the distinction between lying and non-literal speech.