

# Lying and Asserting

Andreas Stokke  
andreas.stokke@gmail.com  
CSMN, University of Oslo

March 2011

— published in *Journal of Philosophy*, CX(1), 2013, 33–60 —

## Abstract

The paper argues that the correct definition of lying is that to lie is to assert something one believes to be false, where assertion is understood in terms of the notion of the common ground of a conversation. It is shown that this definition makes the right predictions for a number of cases involving irony, joking, and false implicature. In addition, the proposed account does not assume that intending to deceive is a necessary condition on lying, and hence counts so-called bald-faced lies as lies.

## 1 Introduction

This paper argues for a definition of lying. The definition is that you lie when you *assert* something you believe to be false. This definition is not new, and many philosophers have held it. Yet these proposals differ from each other in relying on different views of assertion. I argue that the account of assertion needed to define lying is the one familiar from the work of Stalnaker (1978), (1998), (2002) according to which to assert that  $p$  is to propose that  $p$  become *common ground*. So my definition of lying is that you lie when you say something you believe to be false and thereby propose that it become common ground.

Traditionally, many philosophers – e.g., Augustine (c 395b), Bok (1978), Williams (2002), Frankfurt (2005) – have defined lying as saying what you believe to be false with the intent to deceive. But as pointed out by Carson (2006), Sorensen (2007), Fallis (2009), and others, this definition fails to do justice to the phenomenon of what Sorensen calls *bald-faced lies*. These are cases in which the speaker says something she believes to be false but does not intend to deceive her audience, and yet intuitively the speaker is lying.

Here is an example that Carson (2006, 290) gives. A student accused of plagiarism is called to the Dean's office. The student knows that the Dean knows that she did in fact plagiarize. But as it is also well known that the Dean will not punish someone who explicitly denies her guilt, the student says,

(1) I didn't plagiarize.

Although the student says something she believes to be false, she does not intend to deceive the Dean. Yet many philosophers have shared the intuition that the student is lying. Hence this is a bald-faced lie.

In light of this, several writers have proposed assertion-based definitions of lying that specify the assertion component so as to count bald-faced lies as lies. In Sections 2 and 3 of this paper I argue against the proposal in Fallis (2009).<sup>1</sup> We will see that the problems with this view suggest that the assertion component must be understood in a particular way with respect to the saying-meaning distinction. In particular, I consider two kinds of cases that have been discussed in the literature: cases involving *irony* and cases involving *false implicature*. In Sections 4 and 5 I spell out my own definition of lying in terms of common ground and show that it handles all of these cases as well as counting bald-faced lies as lies.

## 2 A Definition and its Shortcomings

### 2.1 Fallis's Definition

Fallis (2009) explicitly endorses the assertion-based view of lying. According to this view, lying is defined as follows:

**The Assertion-Based Definition**

S lies to  $X$  if and only if

(A1)  $S$  asserts that  $p$  to  $X$ , and

(A2)  $S$  believes that  $p$  is false.

However, as Fallis acknowledges, the immediate challenge for the proponent of this definition is to specify the underlying notion of assertion. Here is Fallis proposal:

I think that you *assert* something when (a) you say something and (b) you believe that you are in a situation where you should not say things that you believe to be false. More precisely, you assert something when you say something and you believe that Paul Grice's first maxim of *quality* (namely, "Do not say what you believe to be false") is in effect as a *norm of conversation*. (2009, 33)

Fallis (2009, 35) cautions that this proposal may not amount to a full analysis of assertion. I am only concerned with demonstrating that it fails to capture the aspects of assertion that is needed in order to justify the Assertion-Based Definition. The same proviso will apply to my own version of this account of lying.

So the suggestion is that the relevant notion of assertion to be plugged into the Assertion-Based Definition is that to assert that  $p$  is to say that  $p$  while believing that the following norm of conversation is in effect:<sup>2</sup>

---

<sup>1</sup>In another paper Fallis (2010) discusses various proposals for how to define the specific case of deceptive lying, i.e., the type of lying that bald-faced lying is not an instance of. Here I am only interested in the proposal of Fallis (2009) explicitly designed to cover the broad phenomenon of lying, including lying without the intent to deceive.

<sup>2</sup>See Grice (1989, 27).

### **First Maxim of Quality (FMQ)**

Do not say what you believe to be false.

In turn, then, Fallis's (2009, 34) definition of lying is the following:

#### **Fallis's Definition**

S lies to *X* if and only if

- (F1) *S* states that *p* to *X*,
- (F2) *S* believes that FMQ is in effect, and
- (F3) *S* believes that *p* is false.

I will argue below that cases involving a particular kind of irony constitute counterexamples to the right to left direction of this definition. That is, in these cases (F1)-(F3) are satisfied and yet the speaker is not lying.

As is often noted, irony is the classic case in which what is said and what is meant diverge. For instance, in giving examples of "phenomena which are obviously part of what is meant by the speaker but not part of what her linguistic string means", Carston (2002) says,

The textbook case is irony and its standard characterization is that of saying one thing while meaning the opposite. (2002, 15)

Meaning the opposite of what one says, though, is not the only way of being ironic, and we will consider another class of ironic utterances later. Likewise, we will see that, in general, cases in which what is meant departs from what is said present challenges to any version of the Assertion-Based Definition of lying.

## **2.2 Irony and the First Maxim of Quality**

Fallis himself recognizes that irony is a potential source of counterexamples to his definition of lying. Consider one of his main examples. Having been led into a garbage chute in an attempt to escape blasting storm-troopers, Han Solo sneers,

- (2) The garbage chute was a really wonderful idea. What an incredible smell you've discovered!

This is a straightforward example of the classic kind of irony in which what the speaker means is the negation of what she says. In particular, Solo wants to convey that the garbage chute was a *bad* idea.

Although Solo is clearly not lying, there is a strong intuition that he said something he believes to be false, namely that the garbage chute was a wonderful idea. Fallis accepts this intuition:

He is trying to communicate something that he believes to be true (namely, that the garbage chute was a really *bad* idea). But he is certainly *saying* something that he believes to be false. (2009, 53)

This means that Fallis accepts that (F1) and (F3) are satisfied.

He furthermore accepts that “Grice’s first maxim of quality is arguably in effect. If that were true, he would be lying according to my definition. Thus, my definition would be too broad.” (Ibid.) In other words, if Solo believes that FMQ is in effect in this case, then the example is a counterexample to Fallis’s definition. Accordingly, Fallis denies that (F2) is satisfied. By contrast, I will argue that Fallis’s reasons for thinking that (F2) is not satisfied in the case of irony are inadequate.

According to the orthodox, Gricean conception, irony is an example of *conversational implicature*.<sup>3</sup> That is, it is a speech act in which the speaker *flouts* a maxim of conversation, in this case FMQ. Fallis argues that this means that (F2) is not satisfied in the case of irony:

I contend that the norm is not in effect with respect to Solo’s sarcastic comment and that Solo does not believe that it is. [...] I contend that, by flouting this norm of conversation, Solo turns it off. (2009, 53)

But although understanding irony as involving flouting FMQ is in line with the Gricean analysis of this kind of speech act, the claim that when a speaker flouts a maxim, she turns off the maxim is in direct opposition to it.

Here is Grice’s gloss on his example of an ironic utterance:

It is perfectly obvious to [the speaker] A and his audience that what A has said or made as if to say is something he does not believe, and the audience knows that A knows that this is obvious to the audience. So, unless A’s utterance is entirely pointless, A must be trying to get across some other proposition than the one he purports to be putting forward. (1989, 34)

What does the work in this account is the claim that the audience arrive at the intended content because they recognize that this must be the one the speaker intended to convey, despite the fact that it is not what she strictly speaking said.

The reason the audience is able to make this inference is the general assumption that the speaker is being cooperative and in particular is obeying the maxims. This is clear from Grice’s “general pattern for the working out of a conversational implicature”. According to this general scheme, the hearer’s reasoning process are envisioned as proceeding in roughly the following way:

He [the speaker] has said that p; there is no reason to suppose that he is not observing the maxims, or at least the Cooperative Principle; he could not be doing this unless he thought that q; he knows (and knows I know that he knows) that I can see that the supposition that he thinks that q is required; he has done nothing to stop me thinking that q; he intends me to think, or is at least willing to allow me to think, that q, and so he has implicated that q. (1989, 31)

In other words, the intended content is arrived at because the assumption that it was the one intended to be conveyed is required for squaring what the speaker said with the presumption that she is obeying the relevant maxim.

In our case the relevant maxim is FMQ. With respect to our example, then, this reasoning process can be sketched as follows:

- (a) Solo said that the garbage chute was a wonderful idea.

---

<sup>3</sup>See Grice (1989, 34).

- (b) Solo believes that the garbage chute was a bad idea.
- (c) Solo is observing FMQ.
- (d) Unless Solo intended to convey that the garbage chute was a bad idea, he would not be observing FMQ.
- (e) Therefore, Solo intended to convey that the garbage chute was a bad idea.

This means that Solo believes that FMQ is in effect. The reason is that what it is for Solo to speak ironically, on this picture, is for him to expect his audience to undertake the reasoning process sketched in (a)–(e). Consequently, on the standard Gricean account of irony, Solo believes that FMQ is in effect.

Another way to make the point is to note that what is meant by ‘flouting’ a maxim is captured by saying that the maxim in question is *exploited*. To exploit a maxim the speaker needs to assume that it is in effect. That is, she needs to assume that her audience will take her as observing the maxim, since that is the assumption which will lead them to infer that she intends to convey a content different from what she said. Hence, exploiting a maxim requires that it is operative, and that the speaker believes that it is.

### 2.3 Winking

What more precisely does Fallis mean by ‘turning off’ a maxim of conversation? Consider another of his examples. In order to deceive the judge and get Tony acquitted, Silvio testifies in court:

- (3) Tony was home with me at the time of the murder.

Since Silvio knows that Tony was not home with him at the time, this is a clear case of lying (indeed, lying with the intent to deceive). Further, Fallis claims that

Silvio is not lying if he says to Paulie at the club, “Tony was home with me at the time of the murder” and then winks to indicate that he is not to be taken seriously. By winking, Silvio turns off (or “opts out from the operation” of) Grice’s first maxim of quality with respect to this particular statement [...]. (2009, 35)

But although Silvio is clearly not lying in the case where he is winking, this explanation of the case is misguided.

Clearly, in the winking case, Silvio is intending to communicate the opposite of what he is saying. So from a Gricean perspective, the most natural reconstruction of the winking case is exactly parallel to the case of Solo’s utterance. That is, it is a case in which Silvio flouts FMQ in order to communicate the negation of what he said. Here is the reasoning process:

- (a) Silvio said that Tony was home with him at the time of the murder.
- (b) Silvio believes that Tony was not home with him at the time of the murder.
- (c) Silvio is observing FMQ.

- (d) Unless Silvio intended to convey that Tony was not home with him at the time of the murder, he would not be observing FMQ.
- (e) Therefore, Silvio intended to convey that Tony was not home with him at the time of the murder.

Why is Silvio winking? The purpose of the winking is to furnish the audience with premise (b). In the case of Solo's utterance, there is enough contextual evidence to furnish the corresponding premise (that Solo is not keen on the garbage chute). But in this case, there is not, and therefore some other means of indicating that FMQ is being flouted is called for. Winking is one such means. Hence, there is no basis for claiming that winking turns off FMQ. Consequently, Fallis's Definition incorrectly counts Silvio as lying in the winking case.

In general, then, maxims of conversation are not turned off in situations in which speakers intend to convey something different from what they say. Rather, it is precisely because the maxims are *not* turned off, that speakers can rely on the audience to infer their intended meaning.<sup>4</sup>

Whether one wants to classify Silvio's winking utterance as a case of irony is a delicate matter. It is natural to feel that irony (at least the paradigmatic kind) requires a more straight-faced utterance. However, the main point here is that both of these examples show that Fallis's Definition incorrectly counts as lying cases in which the speaker is intuitively not lying although she says something she believes to be false in order to implicate something she believes to be true but intuitively is not lying.

### 3 Saying, Falsely Implicating, and Pretending

We have seen that denying that (F2) is satisfied in the case of irony is not a viable strategy for defending Fallis's Definition. In light of this, a natural reaction is to look for ways of arguing that, in cases of irony, (F1) is violated. This immediately raises the question of the relation between lying and saying. In this section I take up this question.

#### 3.1 Making as if to Say and Bald-Faced Lies

Fallis's Definition is given in terms of a notion of *stating*. As we saw, Fallis accepts that this notion is, at least roughly, equivalent to the intuitive notion of saying according to which Solo says something he believes to be false. However, one response to what I argued above consists in rejecting this claim and to look for a notion of stating or saying according to which Solo did not say that the garbage

---

<sup>4</sup>A final option is to distinguish between different levels of operation of the maxims. Cf. Grice's (1989, 33) comment that in cases like irony "though some maxim is violated at the level of what is said, the hearer is entitled to assume that the maxim, or at least the Cooperative Principle, is observed at the level of what is implicated." It may be possible to work out a definition like Fallis's which specifies that the speaker is required to believe that FMQ is in effect at a particular level, though it is hard to see how precisely to do so while honoring all the data. In particular, it is hard to see how this suggestion can avoid counting false implicature as lying. (See Section 3 of this paper.)

chute was a wonderful idea. If so, Solo will not count as lying. (I will stick to the notion of saying from now on.)

There is support to be found for this reaction in Grice's (1989, 34) own conception of irony. Grice's suggested that, when speaking ironically, a speaker has merely "made as if to say" the proposition expressed by the sentence in question. This line is explicated by Neale (1992), who argues forcefully that Grice's notion of what a speaker said must be understood in relation to the Gricean program of defining *speaker meaning* in terms of a certain kind of communicative intention. According to Neale,

If *U* utters the sentence "Bill is an honest man" ironically, on Grice's account *U* will not have said that Bill is an honest man: *U* will have *made as if to say* that Bill is an honest man. (1992, 523)

I have no quarrel with this interpretation of Grice. The point to be made here is rather that adopting this notion of saying is not open to a proponent of Fallis's Definition. The reason is that if (F1) is understood in terms of this Gricean notion of saying, the definition will rule out bald-faced lies, and hence will not be able to do the work it was originally intended to do.

Neale points out that

it is Grice's view that a statement of the form 'by uttering *x*, *U* said that *p*' entails the corresponding statement of the form 'by uttering *x*, *U* meant that *p*'. (Ibid.)

Famously, Grice analyzed the notion of what a speaker meant in terms of his concept of audience-directed, communicative intentions. On this understanding, a speaker *S* meant that *p* if and only if, roughly, *S* intended her audience to believe that *p* as a result of their recognizing this intention (cf. Neale (1992, 515)). In other words, that *S* said that *p* entails that *S* meant that *p*, and that *S* meant that *p* entails that *S* intended her audience to believe that *p* as a result of their recognizing this intention. Hence, on this notion of saying, that a speaker said that *p* entails that she intended her audience to come to believe that *p* (by recognizing this intention).

So according to this Gricean picture, it is true that Solo did not say that the garbage chute was a wonderful idea in the example of (2). Fallis's Definition would therefore correctly rule out this kind of irony as a case of lying. However, if this notion of saying is adopted in Fallis's Definition, the definition also rules out cases of bald-faced lies. Take the plagiarizing student. The very reason for classifying this case as one of lying without the intent to deceive is that the student does not intend to produce in the Dean the belief that she did not plagiarize. Therefore, (F1) would not be satisfied in the bald-faced lie case, and so bald-faced lies would not count as lies. Hence, this Gricean conception of saying cannot be adopted for Fallis's purposes.

### 3.2 Constraints on Saying

We can extract from the above a general conclusion. If lying is to be defined in terms of saying, then in order to rule in bald-faced lies, it must be formulated in terms of a notion of saying such that a speaker can say that *p* without this entailing either

that she intends her audience to believe the proposition put forward nor that she intends her audience to believe that she herself believes that proposition. Indeed, both speaker and audience may know the proposition to be false, as in the case of the plagiarizing student.

Let us put this schematically. To capture the phenomenon of bald-faced lies, lying must be defined in terms of a notion of saying which satisfies the following conditions (where  $\Rightarrow$  represents entailment, broadly construed):<sup>5</sup>

(C1)  $S$  says that  $p \not\Rightarrow S$  intends her audience to believe that  $p$ .

(C2)  $S$  says that  $p \not\Rightarrow S$  intends her audience to believe that  $S$  believes that  $p$ .

(C3)  $S$  says that  $p \not\Rightarrow S$  believes that  $p$ .

(C4)  $S$  says that  $p \not\Rightarrow S$  knows that  $p$ .

Any definition of lying in terms of a notion of saying not satisfying these conditions will fail to capture the cases we want to capture. That is, (C1-4) constitute constraints on the notion of saying in the sense that if lying is a matter of saying that  $p$  while believing that  $p$  is false (plus some other condition), then the notion of saying must satisfy (C1-4).

We have already seen two candidate notions. On the one hand, defining lying in terms of the strong Gricean notion of saying discussed above violates (C1), because on that notion if  $S$  says that  $p$ ,  $S$  intends her audience to believe that  $p$ . But on the other hand, the intuitive, weak notion of saying according to which Solo said that the garbage chute was a wonderful idea satisfies (C1-4). Given this, a reasonable conjecture is that a satisfactory notion of saying satisfying (C1-4) will be such that ironic speakers of Solo's variety count as performing it.

### 3.3 Falsely Implicating

At this point some might be attracted to a suggestion to define lying in terms of *what is meant* rather than what is said. That is, one would claim that the right definition of lying is that you lie if you convey, or mean, something you believe to be false (plus some other conditions). Assuming that this kind of theory could be spelled out satisfactorily, it would meet the challenge posed above. It would predict that Solo's utterance is not a lie, since what Solo meant – that the garbage chute was a bad idea – is not something he believes to be false.

This suggestion amounts to the claim that, in suitable conditions, falsely implicating is a form of lying. However, this claim is widely rejected in the literature on lying. Indeed, even writers who think that, in certain circumstances, lying and falsely implicating may be equally blameworthy, still take care to distinguish the two phenomena. For example, Adler (1997) argues that while falsely implicating is not lying, it is a form of deception that in some circumstances can be just as morally and epistemically wrong as lying. Adler's motivations arise directly from an assertion-based conception of lying. According to Adler, in the standard case,

<sup>5</sup>(C4) is strictly speaking redundant in that it follows from (C3) since cases in which the speaker does not believe that  $p$  are also cases in which she does not know that  $p$ . I include it for clarity.



lying is a significantly worse choice than other forms of deception. Both choices aim for the victim to believe falsely, but only lying does so through asserting what one believes false. (1997, 435)

Whether or not one ultimately wants to claim, as Adler does, that the impropriety of deception by means of false implicature and of lying is, at least sometimes, equal, few would disagree that the two are to be distinguished.

Adler discusses the following biblical example:<sup>6</sup>

Abraham, venturing into a dangerous land and fearing for his life if Sarah is taken as his wife, tells Abimelech the king that she is his sister. God appears to Abimelech to warn him away from taking Sarah because "She is a married woman." Frightened, Abimelech confronts Abraham, who defends his obvious deception by denying that he lied:

... they will kill me for the sake of my wife. She is in fact my sister, she is my father's daughter though not by the same mother; and she became my wife... (1997, 435)

Most commentators, going back at least as far as Augustine (c 395a), have defended Abraham as not having lied, although he is guilty of deception.

More precisely, the kind of deception Abraham perpetrates is that of implicating something he believes to be false.<sup>7</sup> Suppose that Abraham's original utterance was:

(4) She is my sister.

The obvious, Gricean way to explain this case is as exploiting Grice's First Maxim of Quantity:<sup>8</sup>

**First Maxim of Quantity**

Make your contribution as informative as is required (for the current purpose of the exchange).

The king will take Abraham as having implicated that Sarah is not his wife because that assumption is needed to make his uttering (4) consistent with the presumption that he is obeying the First Maxim of Quantity.

But if lying is defined in the way suggested above, i.e., as conveying what you believe to be false, Abraham counts as having lied, contrary to the judgements of most writers on the subject. Indeed, Fallis agrees that

you are not lying if you make a statement that you believe to be true. In fact, you are not lying even if you intend to deceive someone by making this statement. (2009, 38)

Abraham makes a statement that he (correctly) believes to be true, namely that Sarah is his sister, while intending to deceive the king by implicating that she is not wife. But intuitively, he is not lying.

<sup>6</sup>See *Genesis* 20: 11-13.

<sup>7</sup>Adler (1997, 438) indicates that he does not consider Abraham's utterance as a case of conversational implicature, but rather as of another type of pragmatic inference. By contrast, I think it is plausible to treat the case as conversational implicature, in the way I suggest in the text. For those who do not agree with this line on the example, I refer to the other cases of false (conversational) implicature that Adler discusses. The point I am interested in making concerns this phenomenon, and is independent of what the right analysis of this particular case turns out to be.

<sup>8</sup>See Grice (1989, 26).

Consequently, lying cannot be defined in terms of what is meant, or conveyed, but must be defined in terms of what is said. We saw that the notion of saying that is required is one that satisfies (C1-4), and I suggested that a plausible notion of this kind will be one that counts Solo as saying that the garbage chute was a wonderful idea. So we need another way of not counting Solo's utterance as a lie.

### 3.4 Pretending to Assert

Here is one way of spelling out the general thought behind Fallis's original account: While Solo says something he believes to be false, he is not asserting that proposition. I think this suggestion is basically right; although we have seen that Fallis's own way of cashing it out, in terms of FMQ, is inadequate. What is right about the proposal is the idea that to assert that  $p$ , one must say that  $p$ , although one can say that  $p$  without asserting that  $p$ . The bald-faced liar asserts what she says, which is something she believes to be false, and hence she is lying. But we need a way of denying that Solo asserts what he says, so as to avoid calling him a liar.

The common ground definition that I will offer in the next two sections preserves this basic strategy. That is, I will agree that asserting entails saying, though the converse is not the case. Before turning to that, though, I want to look at one other way of pursuing this strategy, because doing so is convenient for bringing out a further type of case that we need to take into account.

According to writers like Clark and Gerrig (1984), Recanati (2004), and Currie (2010), in speaking ironically a speaker merely *pretends* to assert. For example, Currie says that

as exploiters of irony we engage in pretence. We pretend to congratulate, approve, admire, and, occasionally, to criticize and deplore. (2010, 150)

If one follows these authors in maintaining that irony involves merely pretending to assert, then Solo's utterance does not count as a lie because Solo is not asserting.<sup>9</sup>

Note that this strategy avoids the problem we articulated for the strong, Gricean notion of saying. That is, it does not end up denying that bald-faced lies are lies. Assertion theorists of lying are usually adamant that bald-faced lies are not merely pretended assertions. A clear case is Sorensen (2007). Sorensen's main example is of an Iraqi minder, Takhlef, to a visiting journalist. During the visit, the journalist realizes that Takhlef in reality is not sympathetic to the regime. Yet, Takhlef explicitly affirms that

- (5) Everything President Saddam Hussein did in the past was good and everything he will do in the future is good.<sup>10</sup>

Since Takhlef believes that (5) is false, and has no illusions as to the convictions of the journalist, this is a case of bald-faced lying. Sorensen says about this example,

<sup>9</sup>Recanati (2004, 19) and Camp (forthcoming, 6–7) note that the Gricean proposal according to which ironic speakers merely make as if to say is naturally understood as a version of this line. Yet, of course, this does not mean that any version of the pretense view of irony is committed to the strong Gricean notion of saying which violates (C1-4). See also Soames (2008) for another view on which neither what is said nor what is implicated count as asserted in cases like Solo's utterance.

<sup>10</sup>Sorensen takes this example from Seierstad (2003).

Takhlef is not merely pretending to assert that Saddam's leadership is perfect. He wants to be on the record. He defends the proposition by words and deeds. (2007, 252)

Similarly, Carson (2006) comments on the case of the plagiarizing student:

The student says this on the record in an official proceeding and thereby warrants the truth of statements he knows to be false. (2006, 290)

In other words, the claim would be that although someone speaking ironically merely pretends to assert, the bald-faced liar is engaged in genuine assertion. Hence, the latter is lying, but the former is not.

One difficulty with this strategy is that, as pointed out by, e.g., Camp (forthcoming), there are cases of irony in which the speaker is intuitively engaged in wholesale assertion. As Camp says,

irony is compatible with the speaker's genuinely committing herself to some content by her utterance. (forthcoming, 11)

Consider an example discussed by Camp:

- (6) The hotel room costs a thousand dollars a night. Of course, for that you get a half bottle of Australian champagne *and* your breakfast thrown in.<sup>11</sup>

Camp's verdict on the case is that the irony it involves "targets just the implicature that the room's apparently high expense is significantly offset by the half bottle of Australian champagne and breakfast; the sentence meaning is itself presented 'straight'." (forthcoming, 9)

I agree that in this case the speaker is genuinely asserting what she says. In fact, an important piece of evidence for this is that there are cases of this kind in which the speaker is lying. Suppose for example that the speaker is well aware that you actually get two dozen bottles of French champagne, free 24 hour room service, a private chef, and a personal butler with the room. In such a situation, an utterance of (6) is clearly a lie. And a natural way of explaining this is to say that the speaker has asserted something she believes to be false. Namely that you get half a bottle of Australian champagne and breakfast. (Of course, in this case, the implicature is also false; but, as we have already seen, falsely implicating is not lying.)

So it cannot be argued that *in general* irony involves pretended assertion, and hence irony *per se* cannot be analyzed in terms of pretense. To be sure, this does not rule out the possibility that the difference between the two kinds of irony that we have considered is precisely the presence or absence of pretense, although the underlying phenomenon that they both exemplify is to be analyzed with reference to some other category. If one wants to put one's money on this, one can point to this difference in order to count Solo as not lying while counting the hotel room guest as lying in the scenario described.

As I suggested above, given an analysis of saying that satisfies (C1-4), this line could potentially provide a way of employing the claim that while all assertions are sayings, not all sayings are assertions in order to arrive at a satisfactory definition of lying. And as we will see, the common ground definition can be taken as an

---

<sup>11</sup>Example from Bredin (1997).

example of this. I am not in principle opposed to the claim that Solo is pretending to assert. But, as I will suggest, the proponent of the common ground view does not need to commit to the claim that the difference between the two kinds of ironic utterance is to be spelled out in terms of pretense. That is, the common ground view offers a less costly explanation of why Solo is not asserting, although he is saying something.

## 4 The Common Ground Definition of Lying

In this section and the next I spell out my definition of lying and show that it makes the right predictions. I begin in this section by introducing the conception of assertion it relies on, and I then demonstrate how we handle the cases discussed in the paper so far. In the next section I then turn to some further types of examples and I discuss some potential problems for the view I favor.

### 4.1 Assertion and Common Ground

I propose to define lying in terms of the account of assertion developed by Stalnaker (1978), (1998), (2002). In doing so we are not committing ourselves to a view of the final analysis of the phenomenon of assertion.<sup>12</sup> As with Fallis's account, what we are after is just to capture the aspects of assertion that are relevant for defining lying.

According to Stalnaker's influential theory, conversations evolve against a background of mutually shared information called the *common ground*. An assertion is a proposal to add to, or update, the common ground with new information. But we need to be careful in spelling out some of the details of this theory. First, it will be important to specify, as we did earlier, that asserting essentially involves saying. Let me explain why.

The common ground is used for many purposes, such as keeping track of pre-suppositions and evaluating indexicals. For example, for an utterance of (7a) to be felicitous, it needs to be common ground that there is a president, and for an utterance of (7b) to be felicitous, it needs to be common ground who the speaker is.

- (7) a. The President is in Washington.
- b. I bought a new car.

Correspondingly, there are many ways in which a piece of information may become common ground. For example, it may be common ground that there is a president before an utterance of (7a), just because this is something that people know or are expected to know. But an utterance of (7a) may also have the effect of making it common ground that there is a president through the familiar mechanism that Lewis (1979) called *accommodation*. Similarly, Stalnaker (1998, 101) points out that a use of *I* does not require that it be common ground who the speaker is before the

---

<sup>12</sup>For recent discussion of this view of assertion, see e.g., Egan (2007), Hawthorne and Magidor (2009), Stalnaker (2009), MacFarlane (2010).

utterance is made; that information may become common ground simply because it is obvious who is making the utterance.<sup>13</sup>

The Stalnakerian view, as I will understand it, is that an assertion is a proposal to update the common ground with something that is said. So to assert that *p* is to say that *p* and thereby propose that *p* become common ground.<sup>14</sup>

Next, it will be important to be precise about how the notion of common ground information is to be understood. According to one version of this framework, the common ground of a conversation is a set of mutually shared beliefs:

It is common ground that *p* in a group if all members believe that *p*, and all believe that all believe that *p*, and all believe that all believe that all believe that *p*, etc.

Adopting this view of the common ground means that asserting that *p* is a proposal for the participants (including the speaker) to believe that *p*. Most commonly, therefore, the speaker herself will already believe that *p* when she asserts it. By asserting that *p* the speaker is inviting the other participants to share her belief that *p*.<sup>15</sup>

Obviously, this version of the common ground view of assertion will not do for our purposes. If to assert that *p* is to say that *p* and thereby propose that *p* become common ground, and the common ground is defined in terms of belief, then the notion of saying will violate (C1-4). For example, the student in the Dean's office is not proposing that they both believe that she did not plagiarize.

Fortunately, as I explain below, this version of the common ground view of assertion is not only inadequate for our purposes, but also does not do justice to the way the account is envisioned by Stalnaker. Instead, as we will see, the correct version is just what we need to define lying in terms of assertion.

Stalnaker (2002) has emphasized that the common ground is to be defined in terms of an attitude weaker than belief. The main reason is that common ground information that is known (or believed) to be false is no obstacle to conversational smoothness.

Successful communication is compatible with presuppositions that are recognized to be false, but the information that they are being presupposed must be actually available, and not just assumed or pretended to be available. (2002, 716)

Instead of belief, then, Stalnaker proposes to define common ground information in terms of *acceptance*. Acceptance is a non-factive propositional attitude weaker than

---

<sup>13</sup>Another way for information to become common ground is by perceptual salience. Stalnaker (1978, 86) points out that if a goat walks into the room, it may become common ground that a goat is in the room, as witnessed by the fact that one can felicitously refer to the goat in a question like, 'How did that thing get in here?' I am not concerned with every alternative here. See Stalnaker (1978), (1998), (2002) for discussion. Sophisticated theories of discourse salience in this tradition are found in, e.g., Heim (1982), Roberts (2002), (2003).

<sup>14</sup>This leaves open the possibility that not only strictly speaking truth-conditional information can be the object of assertion, in this sense, as long as it is part of what is said. The relation between what is said and semantic content is debated in the semantics-pragmatics literature. For views on the pragmatic end of this scale, see, e.g., Bach (1994), Carston (2002), Recanati (2004). For opposition to these views, see, e.g., Stanley (2000), Cappelen and Lepore (2004), Predelli (2005), Stokke (2010).

<sup>15</sup>For a version of this conception of the common ground view of assertion, see Egan (2007).

belief. That is, that a subject  $S$  accepts that  $p$  does not entail that  $p$  is true, nor that  $S$  believes that  $p$ .

However, it would be inadequate to simply define the common ground as a set of propositions that the participants mutually accept. Interestingly, Stalnaker uses the example of lying in order to motivate this point:

Even the liar, if he really intends to communicate, has to believe that the information needed to interpret his lies will really be common ground. So we might identify the common ground with common *belief* about what is accepted. (2002, 716)

Although this remark is not developed further by Stalnaker, I take it to be obvious that what is meant by this is the following. Suppose that someone wanted to lie by uttering (7a). For example, suppose the president is not in Washington and the speaker knows this, but wants to deceive his hearers into believing that she has. Then Stalnaker's point is that, just as in the normal situation, the speaker has to rely on it being common ground that there is a president. Lies exploit the common ground in the same way as sincere assertions.

In light of the fact that common ground information is not required to be the object of full scale belief, although it must be believed to be available, Stalnaker proposes the following definition of the common ground (Ibid.):

It is common ground that  $p$  in a group if all members accept (for the purpose of the conversation) that  $p$ , and all believe that all accept that  $p$ , and all believe that all believe that all accept that  $p$ , etc.

So the view of assertion that I want to adopt for the purpose of defining lying is the following: To assert that  $p$  is to say that  $p$  and thereby propose that  $p$  become common ground, where common ground is understood in the way just described.<sup>16</sup>

## 4.2 The Common Ground Definition

Here is my definition of lying:

### The Common Ground Definition

$S$  lies to  $X$  if and only if

- (L1)  $S$  says that  $p$  to  $X$ , and
- (L2)  $S$  proposes that  $p$  become common ground, and
- (L3)  $S$  believes that  $p$  is false.

There are two crucial notions involved in this definition, namely the notion of saying in (L1) and the notion of common ground in (L2). In order to capture the phenomenon of lying, both of these notions need to be understood in particular ways.

I have already introduced these qualifications above. The notion of common ground mentioned in (L2) is to be understood as defined in terms of belief about

---

<sup>16</sup>It is not ruled out that some common ground information is the object of outright belief, in addition to acceptance. Obviously, very often, the information we rely on, that is, accept, is genuinely believed, if not known. All we are assuming here is that the necessary condition on common ground information is that it be accepted.

what is accepted for the purpose of the exchange. Similarly, the notion of saying in (L1) is to be understood such that it satisfies the constraints (C1–4) articulated earlier. In addition, in accordance with how I have used it in this paper, the notion of saying in this clause is intended to rule out what I have referred to as what is meant by a speaker, e.g., conversational implicatures.

These assumptions are in need of elaboration and more specifications are required with respect to both notions. I explain these points in the rest of this section and the next by going through how the Common Ground Definition handles particular cases.

### 4.3 Irony and False Implicature

The Common Ground Definition counts Solo as not lying. The reason is that while Solo says something he believes to be false – that the garbage chute was a wonderful idea – he is not proposing to make that proposition common ground. This is summarized here:

#### **Solo's Utterance**

(L1) satisfied: Solo says that the garbage chute was a wonderful idea.

(L2) not satisfied: Solo does not propose that it become common ground that the garbage chute was a wonderful idea.

(L3) satisfied: Solo believes that the garbage chute was a bad idea.

#### **Prediction: Solo is not lying.**

On the other hand, the Common Ground Definition counts the hotel guest as lying, since she does propose to make common ground a proposition that she says while believing it to be false:

#### **The Hotel Guest**

(L1) satisfied: The hotel guest says that you get half a bottle of Australian champagne and breakfast.

(L2) satisfied: The hotel guest proposes that it become common ground that you get half a bottle of Australian champagne and breakfast.

(L3) satisfied: The hotel guest believes that you get two dozen bottles of French champagne, free 24 hour room service, a private chef, and a personal butler.

#### **Prediction: The hotel guest is lying.**

So the Common Ground Definition correctly handles both types of irony.

This account, then, provides a way of explaining the difference between these two kinds of irony. I suggested earlier that the difference is that in the classic form

of irony the speaker is not asserting what is said, whereas in cases like that of the hotel guest the speaker is asserting what is said. This difference is explained by observing that whereas Solo is not proposing to update the common ground with what is said, the hotel guest is. I take this to be obvious: Solo is not proposing that the information that the garbage chute was a wonderful idea become part of what is commonly accepted for the purpose of the conversation. But the hotel guest does propose that the false information that you get half a bottle of Australian champagne and breakfast become part of what is accepted. So the Common Ground Definition not only predicts the correct result that Solo is not lying, while the hotel guest is; it does so by offering a plausible explanation of the intuitive difference between the two forms of irony.

As we saw, another option is to explain the difference by claiming that Solo merely pretends to assert, while the hotel guest is genuinely asserting. If one accepts the common ground view, one has a way of explaining the difference without having to endorse this claim about pretense. Of course, one is not precluded from also accepting that Solo is pretending to assert, if one finds that claim appealing. But, as advertised, the proponent of the common ground view can remain neutral on this issue, since she has a way of marking the difference: that Solo is not asserting what is said is explained by recognizing that what is said is not proposed for common ground inclusion.

Let us now turn to cases of false implicature. In our example Abraham says that Sarah is his sister and thereby falsely implicates that Sarah is not his wife. On the Common Ground Definition, Abraham counts as not lying, because while he proposes to make what he said common ground, he also believes that what he said is true:

#### **Abraham's Utterance**

(L1) satisfied: Abraham says that Sarah is his sister.

(L2) satisfied: Abraham proposes that it become common ground that Sarah is his sister.

(L3) not satisfied: Abraham believes that Sarah is his sister.

#### **Prediction: Abraham is not lying.**

So the definition correctly avoids counting false implicatures as lies.

The results of the Common Ground Definition for cases involving irony and false implicature that I have just summarized rest on the assumption that there are differences between cases in which saying and meaning depart depending on whether what is said is proposed for inclusion in the common ground. The cases where what is said is proposed for inclusion in the common ground are cases in which the speaker intuitively asserts what is said. Such are the case of the hotel guest and of Abraham's false implicature. Given the conception of lying advocated here, this kind of case is compatible with lying. By contrast, cases in which the speaker intuitively does not assert what is said – such as Solo's utterance – are not compatible with lying.



Someone might object that since, in all of these cases, what is meant is proposed for common ground inclusion, the common ground conception of assertion is committed to the claim that what is meant is asserted. This would be problematic. As Soames (2008) rightly says regarding Grice's (1989, 33) well-known letter of recommendation case,

the proposition implicated – that the job candidate is no good – is the real point of the writer's remark. Although this may tempt one to identify the implicature as the writer's "real assertion," the temptation should be resisted – since the whole purpose of using indirect means to convey this information was to avoid having to state it. (2008, 443)

However, the common ground conception of assertion that I am relying on does not count implicatures as asserted, because although implicatures are proposed for the common ground, they are not part of what is said. To repeat, according to that conception, a speaker asserts that  $p$  only if she says that  $p$  and thereby proposes to update the common ground with  $p$ . Asserting requires saying, and in turn, so does lying.

The Common Ground Definition makes the right predictions in all the cases involving the saying-meaning distinction that we have considered. In contrast to Fallis's definition, it does not count as lies standard cases of irony. And it meets the challenge of doing so while counting other kinds of irony, exemplified by the hotel guest, as compatible with lying. Finally, since the definition is couched in terms of what is said, it does not fall into the trap of counting false implicatures as lies.

#### 4.4 Bald-Faced Lies

What does the Common Ground Definition say about bald-faced lies? Consider the plagiarizing student. Here is what we predict for this case:

##### **The Plagiarizing Student**

(L1) satisfied: The student says that she did not plagiarize.

(L2) satisfied: The student proposes that it become common ground that she did not plagiarize.

(L3) satisfied: The student believes that she did plagiarize.

##### **Prediction: The student is lying.**

So the Common Ground Definition correctly rules in bald-faced lies.

The central point to note here is the claim that (L2) is satisfied in the cases of bald-faced lies, that is, the suggestion that the bald-faced liar is proposing to add what is said to the common ground. The shared intuition about the case of the plagiarizing student is that the reason the student makes her utterance – despite the fact that both she and the Dean know full well that what she says is false – is that she wants to be on the record in order to be sure to avoid punishment. This is plausibly explained in terms of the common ground. Namely, to say that the student wants to be on the record is just to say that the student wants it to be common ground (in the weak sense we are assuming) that she did not plagiarize.

To ratify this proposal, there are two main issues that need to be clarified. Namely, should we assume that the information that the student is guilty is common ground from the outset, and if so, is that a problem for our account of lying?

First, it might be suggested that, given that the Dean and the student both know that the student plagiarized, that information is common ground between them when the conversation in the Dean's office begins. Indeed, we can assume for the sake of argument (as is plausible) that it is *common knowledge* that the student plagiarized, i.e., that each of them knows that she did, and each of them knows that the other knows that she did, etc.

However, the assumption that if a proposition  $p$  is common knowledge, then  $p$  is common ground is highly implausible. The main reason is that a proposition can be common ground even if its negation is common knowledge. That would be impossible if common knowledge automatically became common ground, since the common ground of a conversation cannot contain both a proposition and its negation. For example, consider Donnellan's (1966) classic example. At a cocktail party, Alice says to Bob,

(8) The man drinking a martini is a philosopher.

As Stalnaker (2002) points out, successful communication in this case does not depend on mutual knowledge (or even belief) that the man is drinking a martini:

perhaps it is mutually recognized that it is not a martini, but mutually recognized that both parties are *accepting* that it is a martini. The pretense will be rational if accepting the false presupposition is an efficient way to communicate something true – information about the man who is falsely presupposed to be the man drinking a martini. (2002, 718)

In this case, it is common knowledge (let us assume) that it is not a martini and yet it is common ground that it is a martini.

It cannot be assumed, therefore, that simply because it is common knowledge between the Dean and the student that the student plagiarized, that information is common ground. However, there might of course be particular circumstances about the case due to which, it is in fact common ground that the student did plagiarize.

This brings us to the second point. For even if we assume (although, I think, implausibly) that it is initially common ground that the student plagiarized, this does not threaten the prediction of the Common Ground Definition that the student is lying when she makes her utterance. The reason is that it is nevertheless clear that the student is proposing to update the common ground with her utterance. The reason she says what she says is to make sure that the common ground comes to include the false information that she did not plagiarize. The student wants herself and the Dean to mutually accept that she did not plagiarize. And this may come about as the result of altering previous common ground information, or of updating with new information.

## 5 A Refinement

We have seen that the Common Ground Definition makes the right predictions for bald-faced lies, both types of irony, and false implicatures. I take this to be a strong

recommendation of the theory. But as we will see in this section, given some plausible assumptions, the definition is also able to handle a range of more complicated cases.

## 5.1 The Confused Politician

Here is an example discussed by Carson (2006) and Fallis (2009). A politician is invited to give a humorous speech at a festive banquet and a serious speech at a formal banquet. She confuses the dates, and ends up delivering the humorous speech at the formal banquet and the formal speech at the festive banquet.<sup>17</sup>

Take the first event. Suppose that during the humorous speech the politician tells a story about the President having “broken wind” during a meeting with some ambassadors. The politician knows that this event did not actually happen and is only relating it to make a joke. The common verdict on this case is that the politician is not lying. She was only joking, although her audience were expecting something else.

Now consider the second event. Suppose that during the speech she says something she knows to be false, say that the President withheld important information. Fallis’s (2009) intuition, which I agree with, is that the politician is lying in this case. She is speaking in earnest, and she is therefore lying when she says something she believes to be false, even though the audience was expecting her to be informal.

How does the Common Ground Definition handle this case? In order to handle the cases, we need to make a qualification about common ground information. I spell out this qualification below. Once it is in place, here is what we will predict for these cases:

### **Confused Politician 1: Humorous Speech**

(L1) satisfied: The politician says that the President broke wind.

(L2) not satisfied: The politician does not propose that it become common ground that the President broke wind.

(L3) satisfied: The politician believes that the President did not break wind.

### **Prediction: The politician is not lying.**

So what I will suggest below is that the reason the politician is not lying in the case where she delivers the humorous speech at the formal banquet is that she is not proposing to update the common ground (in the relevant sense) with what is said. By contrast, then, we will make the following prediction about the converse situation:

### **Confused Politician 2: Serious Speech**

---

<sup>17</sup>Carson (2006) used the former situation to motivate an aspect of his own definition of lying. Fallis (2009) points out that, even so, Carson still makes the wrong prediction about the latter situation. I agree with Fallis’s on this point, and therefore consider Carson’s theory to be out of the running, at least until further developed. See Fallis (2009, 47–48) for details.

(L1) satisfied: The politician says that the President withheld information.

(L2) satisfied: The politician proposes that it become common ground that the President withheld information.

(L3) satisfied: The politician believes that the President did not withhold information.

**Prediction: The politician is lying.**

Let me now turn to the motivation for these claims.

## 5.2 Official and Unofficial Common Ground

The motivation for claiming that (L2) is not satisfied in the case where the politician is joking, whereas (L2) is satisfied in the case where she is speaking in earnest, is based on a distinction between what I shall call *official* and *unofficial* common ground. It is convenient to introduce this distinction by considering another potential source of problems for the Common Ground Definition, namely cases of *assumption*.

It is often noted that the weak definition of common ground that we have relied on, i.e., the conception of common ground information in terms of acceptance for the purpose of the exchange, allows mere assumptions to become common ground. Indeed, Stalnaker remarks that

One may make *assumptions*, and what is assumed may become part of the common ground, temporarily. (2002, 704)

This kind of case is a potential threat to the Common Ground Definition of lying. The definition is in danger of over-generating in that it is obvious that one is not lying if one makes an assumption that one believes to be false, e.g., for the sake of reductio, or conditional proof, or the like. However, there are two reasons for thinking that the Common Ground Definition is not ultimately falsified by these cases.

First, it is not obvious that there are clear examples of assumption for the sake of argument in which the speaker *says* the false proposition, in our sense of this term. Concretely, what is the utterance that one would make? One candidate utterance involves a prefix of the form 'Suppose/assume that...'. For example, I might begin a reductio as follows:

(9) Assume that there is a set of all sets that are not members of themselves.

It is clear that, although I believe that there is no such set, I am not lying in uttering (9). But have I said that there is a set of all sets that are not members of themselves? It seems most natural to think that I have not. (9) is an imperatival, not a declarative, sentence and correspondingly is most naturally taken as issuing a request or instruction rather than directly expressing a proposition.<sup>18</sup> Very few theories are going to consider this kind of utterance candidate for lying, then.

<sup>18</sup>For relevant discussion, see Austin (1962, 86–89), Searle (1969, 161–162), Green (2000).

But there are other utterances I can make. For example, I can start the proof like this:

- (10) Let's make an assumption. There is a set of all sets that are not members of themselves.

However, even though the second sentence is in the declarative mood, it is not evident that I count as having said that there is a set of all sets that are not members of themselves. Indeed, it is difficult to pin down the relevant difference between this case and that of (9). There is an intuitive sense in which I have not asserted the existence of the notorious set even in the case of (10).

It is at least not obvious, then, that the proponent of the Common Ground Definition of lying will be forced to concede that (L1) is satisfied in cases of assuming or supposing something believed to be false.

But there is a second point to be made here. Namely, even if it can be argued that (L1) is satisfied in these cases, it may still be argued that (L2) is not satisfied. That is, one can question the claim that information merely assumed for the sake of argument are added to the common ground, in the relevant sense.

As we saw, Stalnaker explicitly qualifies his endorsement of this claim by indicating that assumptions are often added with a merely *temporary* purpose in mind. This suggests the following picture. When an assumption is made, a temporary common ground is opened up in which the information used for the purpose of the argument is stored. I propose to call such temporary common grounds 'unofficial' in order to distinguish them from more permanent, 'official', common grounds.

There are good reasons to think that we need to acknowledge the existence of unofficial common grounds for independent reasons. Consider for instance a play being performed. One plausible thought is that during the performance of the play, a common ground is active, which stores information about the characters and situations that the play is about. This unofficial common ground is used to keep track of presuppositions, determine indexical content, etc., expressed by the characters in the play.

Here is one way of motivating such a picture. Imagine we are witnessing the performance of a play set in a monarchy. We are in the middle of an intense scene, where, after the exit of her husband, the hero is declaring his love for the heroine. Indexicals are being used in lines such as, "I love you". And presuppositions are being invoked by utterances involving "the King". Without even thinking about it, we are taking these indexicals to refer to the characters in the play, which suggests that the common ground we are using to evaluate these utterances is one that corresponds to the reality of the play. And similarly, we are resolving presuppositions by using this common ground of the play.<sup>19</sup>

---

<sup>19</sup>Note that this view is compatible with more than one analysis of fictional discourse. According to Lewis (1978), sentences of fiction contain a hidden intensional operator 'In fiction  $f$ ...' defined as follows: "a prefixed sentence "In fiction  $f$ ,  $\phi$ " is true (or, as we shall also say,  $\phi$  is true in the fiction  $f$ ) iff  $\phi$  is true at every possible world in a certain set, this set being somehow determined by the fiction  $f$ ." (1978, 39) Familiarly, the common ground of a conversation is taken to determine a set of possible worlds  $w$  s.t. for all propositions  $p$  in the common ground,  $p$  is true at  $w$ . It is then straightforward to take the set of worlds determined by  $f$  to be what I have called the unofficial common ground. Similarly, the view I am

Suppose now that in the middle of this scene, the cuckolded husband suddenly bursts in, runs to the edge of the stage and shouts in a distressed tone of voice,

- (11) Ladies and gentlemen! I have just received word that the President has been shot!

Most likely, we will all take this utterance to express the proposition that the actor has just received word that the (real) President has been shot. In particular, we will not be puzzled by the fact that there is no president in the play. And we will take *I* to refer to the actor, not the character he was playing a few minutes ago.

A plausible way of modeling what happens in cases like that of the interrupted play, then, is in terms of two common grounds being operative at the same time. Of one of these common grounds we know that it is the official one, namely the one that contains information such as that there is a president, that the person on the stage is an actor, and not a jealous husband from a distant era, etc. (Yet even the official common ground may contain information that everyone knows to be actually false.) It is likely, therefore, that the complex ways in which we smoothly move from discourse to discourse is to be explained by interactions with different common grounds, some more official than others.

### 5.3 Lying and Official Common Ground

The notion of common ground in our definition of lying is intended to be that of the official common ground, i.e., it is intended to rule out cases in which information is stored for temporary purposes such as play acting, etc. Indeed, notice that the actor is clearly lying if he believes that it is false that the President has been shot, when he bursts in and makes his announcement. The reason for this, I am suggesting, is that in this case he is proposing to update the official common ground.

Likewise, we can now explain the assumption we made about the cases involving the confused politician. In the case where she is delivering her humorous speech, the politician is not proposing to add the information that the President broke wind to the official common ground, but merely to the unofficial common ground of her joking speech. Hence, she does not count as lying. By contrast, when she delivers her serious speech, she is intending to update the official common ground, and she therefore counts as lying when she says something she believes to be false.

Note that it does not matter that the audience, in each case, is expecting the opposite mode of speech. Our definition of lying merely requires that the politician is *proposing* to add to the (official) common ground. In general, lying does not require that the information which the speaker believes to be false be *in fact* included in the common ground; even if the other participants refuse to accept the speaker's false information, she still counts as lying just for proposing to update with something she knows to be false.

---

sketching is compatible with accounts such as the one proposed by Predelli (2005, 66–73) according to which the hallmark of sentences of fiction is that they are evaluated at fictional worlds, although this shift is not triggered by a hidden operator. See Recanati (2000, ch. 15) for an alternative view of utterances inside fictions.

So 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. The same mechanism accounts for the case of Silvio's winking utterance. Although there is a sense in which Silvio is updating an unofficial common ground, he does not count as lying because he is clearly not intending to update the official common ground, as signaled by his winking.

What about Solo's utterance? Does a speaker of this kind of ironic utterance propose to update an unofficial common ground with what is said? The answer to this question will depend on whether it is found that information delivered as what is said ironically in this type of utterance may for some purposes be treated as common ground in the sense of being used, for instance, for supporting presuppositions or interpreting indexicals.

A full investigation of this matter lies far beyond the scope of this paper. But it is clear that, even if it is found that what is said by ironic speakers may sometimes be used as common ground information further downstream in the conversation, this will not be a threat to our claim that Solo is not asserting (and hence is not lying), since we can reasonably maintain that the relevant kind of common ground is unofficial. Indeed, it is plausible that *if* what is said ironically can be used later on in this way, doing so will involve a sense of playfulness or simulation, as if someone should reply to a question of who has hairy feet with "Hobbits have hairy feet."<sup>20</sup>

Finally, I want to consider a potential worry concerning what I have just said in relation to bald-faced lying. Someone might object that, given the distinction between the official and the unofficial common ground, the bald-face liar should be taken as proposing to update a merely unofficial common ground, and hence will not be counted as lying after all.

However, this suggestion is clearly not what we are committed to. As emphasized repeatedly, bald-faced lying is characterized by the speaker's desire to be on the record. The plagiarizing student is not joking, she is not play acting, nor is her speech act comparable to mere assumption. She is sincerely asserting her innocence. So it seems natural to say that she is putting forth her statement for the official common ground, and not for some merely temporary purpose.

This suggestion can be corroborated in the following way. Both assertions and presuppositions added to an unofficial common ground can later be unproblematically revoked. For example, suppose that, after the politician has given her humorous speech, someone charges her with having lied. She can defend herself by saying,

(12) No, no, you didn't realize that I was just joking.

And although the politician will be expected to apologize for having made this mistake, she is not obviously reproachable for having lied.

Significantly, the parallel is not the case for the plagiarizing student. If later charged with lying, she cannot claim to have merely been joking, speaking unse-

---

<sup>20</sup>This example also illustrates that an unofficial common ground need not be temporary in the sense of lasting a short time. There are arguably common grounds that we all make use of from time to time, which are unofficial in the sense that we all know that the information they contain is fictional, or the like, but which nevertheless continue to be operative for a very long time.

riously, or the like. In particular, note that even though the plagiarizing student *can* admit later on that she only said what she said in order not to get punished, someone can equally well point out that, even so, she lied.

I think it is safe to conclude that the proponent of the Common Ground Definition<sup>1</sup> has the resources to avoid non-standard discourses like proofs, plays, or jokes being sources of counterexamples.

## 6 Conclusion

You lie when you assert what you believe to be false. Asserting that  $p$  is to say that  $p$  and thereby propose that  $p$  become common ground.

We have seen that this definition of lying successfully handles a variety of examples some of which present counterexamples to other versions of the assertion-based view of lying. The common-ground account of assertion coupled with a suitable notion of saying allows one to predict the intuitively right results in cases involving the saying-meaning distinction, such as the cases of irony and false implicature we have looked at, and in addition to count bald-faced lies as lies.

## References

- Adler, J. (1997). Lying, deceiving, or falsely implicating. *Journal of Philosophy*, 94, 435–452.
- Augustine. (c 395a). Against lying. In R. Deferrari (Ed.), *Treaties on various subjects* (Vol. 16, pp. 125–179). New York: Fathers of the Church, 1952.
- Augustine. (c 395b). Lying. In R. Deferrari (Ed.), *Treaties on various subjects* (Vol. 16, pp. 53–120). New York: Fathers of the Church, 1952.
- Austin, J. (1962). *How to do things with words*. Oxford: Oxford University Press.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Bok, S. (1978). *Lying: Moral choice in private and public life*. New York: Random House.
- Bredin, H. (1997). The semantic structure of verbal irony. *The Journal of Literary Semantics*, 16, 1–20.
- Camp, E. (forthcoming). *Sarcasm, pretense, and the semantics/pragmatics distinction*. Forthcoming in *Nous*. References to online version at: <http://www.sas.upenn.edu/~campe/Papers/Camp.Sarcasm.pdf>.
- Cappelen, H., & Lepore, E. (2004). *Insensitive semantics – a defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
- Carson, T. (2006). The definition of lying. *Nous*, 40, 284–306.
- Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.
- Clark, H., & Gerrig, R. (1984). On the pretense theory of irony. *Journal of Experimental Psychology: General*, 113, 121–126.



- Currie, G. (2010). *Narratives and narrators: A philosophy of stories*. Oxford and New York: Oxford University Press.
- Donnellan, K. (1966). Reference and definite descriptions. *Philosophical Review*, 75, 281–304.
- Egan, A. (2007). Epistemic modals, relativism and assertion. *Philosophical Studies*, 133, 1–22.
- Fallis, D. (2009). What is lying? *Journal of Philosophy*, 106, 29–56.
- Fallis, D. (2010). Lying and deception. *Philosophers' Imprint*, 10, 1–22.
- Frankfurt, H. (2005). *On bullshit*. Princeton, NJ: Princeton University Press.
- Green, M. (2000). The status of supposition. *Nous*, 34, 376–399.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Hawthorne, J., & Magidor, O. (2009). Assertion, context, and epistemic accessibility. *Mind*, 118, 377–397.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases in English*. PhD dissertation, University of Massachusetts at Amherst, 1982. Published by Garland Press, 1988.
- Lewis, D. (1978). Truth in fiction. *American Philosophical Quarterly*, 15, 37–46.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- MacFarlane, J. (2010). What is assertion? In H. Cappelen & J. Brown (Eds.), *Assertion - new philosophical essays* (pp. 79–96). Oxford and New York: Oxford University Press.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15, 509–559.
- Predelli, S. (2005). *Contexts – meaning, truth and the use of language*. Oxford: Clarendon Press.
- Recanati, F. (2000). *Oratio obliqua, oratio recta*. Cambridge, MA: MIT Press.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Roberts, C. (2002). Demonstratives as definites. In K. van Deemter & R. Kibble (Eds.), *Information sharing: Reference and presupposition in language generation and interpretation* (pp. 89–196). Stanford, CA.: CSLI.
- Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26, 287–350.
- Searle, J. (1969). *Speech acts: An essay in the philosophy of language*. Oxford: Oxford University Press.
- Seierstad, A. (2003). *A hundred and one days: A Baghdad journal*. New York: Basic Books. (Trans. I. Christophersen)
- Soames, S. (2008). Drawing the line between meaning and implicature – and relating both to assertion. *Nous*, 42, 440–465.
- Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88, 251–264.
- Stalnaker, R. (1978). Assertion. In *Context and content* (pp. 78–95). Oxford: Oxford University Press, 1999.
- Stalnaker, R. (1998). On the representation of context. In *Context and content* (pp. 96–114). Oxford and New York: Oxford University Press, 1999.

- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Stalnaker, R. (2009). On Hawthorne and Magidor on assertion, context, and epistemic accessibility. *Mind*, 118, 399–409.
- Stanley, J. (2000). Context and logical form. In *Language in context – selected essays* (pp. 30–68). Oxford: Clarendon Press, 2007.
- Stokke, A. (2010). Intention-sensitive semantics. *Synthese*, 175, 383–404.
- Williams, B. (2002). *Truth and truthfulness: An essay in genealogy*. Princeton, NJ: Princeton University Press.