

# Lying and Cooperation

Brian Robinson

## Abstract

This paper argues for a new, cooperation-based definition of lying, according to which a speaker lies to an addressee if and only if the speaker say something false to the addressee and is thereby uncooperative with the purpose of their talk exchange. I review four previously recognized definitional problems, showing that the cooperation-based is able to successfully handle each. Loose talk is introduced as a new definitional hazard. Loose talk is a statement that is strictly false but close enough for conversational purposes. Existing definitions incorrectly find cases of loose talk to be lies. The cooperation-based definition, on the other hand, can explain why saying “It’s 3:30” when the time is actually 3:29 is not a lie.

**Keywords:** lie, lying, cooperation, deception, loose talk, implicature

**Word Count:** 8,859

“No fact is more firmly established than that lying is a necessity of our circumstances.”

— Mark Twain, *On The Decaying Art of Lying*

“‘Lie’ is so unmusical a word.”

— Dowager Countess of Grantham, *Downton Abbey*

## 1 Introduction

Everybody lies. And anybody who says differently is a liar. It seems curious then that defining precisely what a lie is has proven so problematic. Numerous definitions have been proposed, though none has achieved consensus. So far, various authors have presented four definitional hazards that any sufficient theory of lying must navigate:

- **Deception:** whether the speaker must intend to deceive the audience,
- **Addressee:** whether a lie must be addressed to someone, and if so, whom,
- **Implying:** whether one can lie by merely implying something false (or believed false), and

- **Jokes, Irony, Metaphor, and Fiction:** whether a statement can be believed false and nevertheless not be a lie in certain cases.

As assorted alternate definitions have been proposed to pilot these hazards, two main types of definitions have emerged: deception-based definitions and deceptionless definitions. The deception-based definitions maintain that the intent to deceive is a necessary condition for lying. On the other hand, deceptionless definitions reject this condition, opting to characterize lying in terms of what is asserted (or some similar concept) without requiring the intent to deceive.

In this paper, I will add a new, fifth definitional hazard to the list, namely cases of **loose talk**. As it turns out, both deception-based and deceptionless definitions misidentify cases of loose talk as lying. Therefore, I offer a new, cooperation-based definition of lying:

### The Cooperation-Based Definition

Speaker *S* lies to audience *H* if and only if:

- (C1) *S* says *p* to *H*,
- (C2) *S* believes *p* is false,
- (C3) By saying the believed-false *p*, *S* is being uncooperative with the purpose or direction of the talk exchange in which *S* is engaged with *H*.

As a basis for what it means for a speaker to be uncooperative, I am largely borrowing from Paul Grice's Cooperative Principle: "Make your conversational contributions such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" (Grice 1989, pg. 26). I will elaborate upon what it means to lie according to this definition as I review each of the five definitional hazards (and how my cooperation-based definition successfully charts a safe path through them all) in sections 2 through 6. Then in section 7, I will examine in greater detail what it means for a speaker to be uncooperative. For now, it is enough to say that a speaker is uncooperative when his or her utterance does not contribute to the accepted purpose or direction of the talk exchange. Though there are many ways one can be uncooperative with a conversation's purpose, the one we are primarily interested in is by saying something false. When, by saying something believed to be false, one has failed to cooperate with a conversation's purpose, then one has lied. Finally, in section 8, I will briefly conclude with a discussion of how the cooperation-based definition offers the potential to identify the moral wrongness of lying.

## 2 Deception

Liars typically intend to deceive someone by lying. The nature of this deceptive intent, however, requires explication, specifically on three issues: whether deception (and lying) must involve a falsehood, what the intent to deceive is, and whether liars must always intend to deceive.

The first question to answer is whether a proposition  $p$  must be false in order for me to deceive (and lie to) you about  $p$ . Suppose you ask me if there is a department committee meeting today and I say, “No, it’s tomorrow.” But, I’ve gotten the date for the meeting confused; it actually is today. Have I deceived you? Have I lied? It doesn’t seem like I’ve done either. Though my statement may have contributed to you believing something false, I did not intend to produce that result; I was confused. To handle such cases, most definitions of lying focus on the epistemic state of the speaker. To lie and deceive, the liar must *believe* that  $p$  is false (c.f., Augustine c395b, Kupfer 1982, Jones 1986, Adler 1997, Williams 2002). It follows from this condition that one can lie (and intend to deceive) about something true. If  $p$  happens to be true but believed to be false, the liar still intended to deceive.<sup>1</sup>

Next, we have the matter of what precisely is the intent to deceive. I may deceive you by contributing to your acquiring a belief in  $p$ . But that isn’t the only way I can deceive you. Chisholm and Freehan (1977) offer a taxonomy of the intent to deceive. They find eight different ways one can intend to deceive another. For instance, I might allow you to continue to believe  $p$ , contribute to you ceasing to believe not- $p$ , or allow you to continue not believing not- $p$ . Most of our subsequent examples will be those of intending for the audience to acquire a false belief, though I agree that deception is not limited only to them; any of Chisholm and Freehan’s eight types will suffice.

Our third question is much trickier. Is the intent to deceive a necessary condition of lying? The intent to deceive is certainly evident in nearly any case of lying one can readily

---

<sup>1</sup> Carson (2006, pg. 286) also mentions the possibility that the statement be believed not to be true, instead of believed to be false. He admits though to not having any intuitions one way or the other here. As I also don’t have intuitions either way, I have left my definition in the more standard form of the statement being believed to be false, instead of merely not true. Carson also contends that the statement must be false *and* believed by the speaker to be false. Though I don’t share Carson’s intuitions on this point, anyone so inclined could easily amend my cooperation-based definition in this way. Others, such as Bok (1978) and Davidson (1980), exclude any requirement of falsehood; a lie must merely be intended to deceive, even if it is believed to be true.

imagine. In fact, as we'll see later, it is commonly thought that it is a lack of an intent to deceive that keeps jokes, fiction, metaphor, and hyperbole from being lies, despite involving saying something believed to be false. Therefore, some theorists claim that the intent to deceive is a necessary condition of lying (c.f., Chisholm and Freehan 1977, Kupfer 1982, Jones 1986, Frankfurt 1988, Simpson 1992, and Williams 2002).<sup>2</sup> In fact, it is the hallmark of all **deception-based definitions** of lying that they include as a necessary condition that the speaker intend to deceive.

Many theorists, however, have rejected the deception requirement, such as Shibles (1985), Kemp and Sullivan (1993), Carson *et al.* (1982), Carson (1988, 2006), Sorenson (2007), Fallis (2009), and Stokke (2013). Carson (2006) contends that sometimes liars don't expect or intend to deceive with their lies, since they know everyone will see right through the lies, but nevertheless the statements still are lies. He offers the following example that epitomizes the motivation behind rejecting the intent to deceive condition. A student has plagiarized a paper and been reported to the dean. The dean, fearing lawsuits, refuses to penalize any student who flatly denies plagiarizing in an official enquiry. In this case, the dean has seen the evidence and knows with complete certainty that the student has plagiarized. The student knows both these things about the dean. During the enquiry, the student says, "I did not plagiarize." The dean is not fooled, nor did the student expect her to be. There was no intent to deceive. Yet the student still lied, it seems. Carson *et al.* (1982) also offer the example of a criminal pleading not guilty to charges in court. The criminal knows the plea is false (since he knows he committed the crime), but he does not intend or expect the judge or anyone else to believe him. It was merely a false *pro forma* utterance. Yet, many still want to say that the criminal lied. If they are correct, the intent to deceive appears not to be a necessary condition for lying. Sorenson (2007) refers to lies of this sort as bald-faced lies. In instances of bald-faced lies, the speaker and audience believe in common that the statement is false, that each party believes the statement is false, that each party believes that each party believes that the statement is false, and so on. Because the speaker realizes that he or she has these beliefs in common with the audience, there cannot be the intent

---

<sup>2</sup> Frankfurt (1988) and Kupfer (1982) contend that there is a second intention to deceive. A liar, they argue, not only has the deceptive intention that the audience believe the statement to be true, but also that the audience believe that the speaker believes the statement to be true. Chisholm and Freehan (1977) argue that though the intent to deceive is necessary, it need not be the intent to deceive the addressee about the believed-false statement.

or expectation to deceive. While not definitive proof, it is interesting to note that empirical studies have found that the folk generally regard instances of bald-faced lies as lies (Colman and Kay 1981).

If we countenance bald-faced lies as in fact being lies, then an alternate definition of lying is required. Some philosophers (cf., Carson 2006, Sorensen 2007, Fallis 2009, 2010, Stokke 2013) defend what I term **deceptionless definitions**. Two common features unite all definitions in this category. First, they all reject the intent to deceive as a necessary condition. Second, they all stipulate that one lies if one asserts (or something similar) what one believes to be false.<sup>3</sup> Fallis's (2009) definition is a prime example of this category. According to his definition, speaker *S* lies if and only if (1) *S* asserts that *p* to audience *H* and (2) *S* believes that *p* is false. Fallis goes on to elaborate his notion of assertion so that it can avoid identifying instances of irony, jokes, and story telling as lies. As we will see in section 6, these are important test cases for any definition of lying. For now though, it suffices to note that Fallis's definition displays the hallmark of deceptionless definitions in that, without the intent to deceive condition, the notion of assertion (or some similar concept) must perform the definitional heavy lifting in these cases.<sup>4</sup> Though all definitions in this category exhibit this hallmark trait, not all of them explicitly make use of assertion *per se*. Carson (2006), for instance, refers to a speaker "warranting the truth of" a statement.<sup>5</sup> Alternatively, Stokke (2013) refers to a speaker proposing a statement as common ground for the talk exchange. These definitions, though different, are sufficiently close to assertion to be grouped in the deceptionless category. Stokke relies on what it means to propose a statement as common ground to perform the definitional heavy lifting of avoiding the

---

<sup>3</sup> Though not a necessary feature of deceptionless definitions, they typically require that a lie only be believed to be false, instead of actually be false.

<sup>4</sup> Carson (2006, pg. 300) notes that some notions of assertion require that a speaker intend the audience to believe the proposition expressed. As Fallis (2010) correctly points out, such notions of assertion will not work for any deceptionless definition of lying. Fallis (2009, 2010) therefore relies on an alternate notion of assertion. Carson (2006) and Stokke (2013) opt instead for something close to but different from assertion to avoid this confusion. Just because one has proposed *p* for common ground, for instance, does not entail that one intends the audience to believe that *p*.

<sup>5</sup> Carson (2006, pg. 299-300) explicitly argues against claims that lying is asserting something one believes false. Instead, he argues that his concept of warranting the truth more accurately defines lying. Nevertheless, his definition is sufficiently structurally similar to others in the deceptionless category to be grouped with them.

troublesome cases of irony, fiction, etc.; Carson does the same for warranting the truth.<sup>6</sup>

Returning to the matter at hand of bald-faced lies, deceptionless definitions claim that bald-faced liars asserted (warranted, proposed for common ground) something they believed to be false. So they still lied, even without the intent to deceive. The intent to deceive, though common, isn't necessary for lying, according to these definitions.

For my part, I agree that generally bald-faced lies are lies, and so the intent to deceive is not a necessary condition for lying. As such, the cooperation-based definition makes no mention of deception. According to my definition, bald-faced lies are lies because the speaker is being uncooperative. To see why, we need to review Grice's notion of cooperation in talk exchanges. Grice notes that speakers typically don't make random, disconnected remarks. Rather they tend to be rational and therefore cooperate toward a common purpose or direction when they speak. To observe the Cooperative Principle, Grice claims, speakers will follow (or ostentatiously flout) his conversational maxims grouped under the four categories of Quality, Quantity, Manner, and Relation. Observation of these maxims is supposed to be how we (typically) cooperate with the purpose of the talk exchange.<sup>7</sup> For instance, the maxim of Relation requires that a speaker's comments be relevant to that purpose. The first maxim of Quality stipulates, "Do not say what you believe to be false" (1989, pg. 27). Audiences are supposed to assume that speakers are observing the Cooperative Principle and the attendant maxims in order to work out what speakers have conversationally implicated (Grice 1989, pg. 26; Neale 1992).

Whether an apparent bald-faced liar is being uncooperative is then a matter of interpretation. We must determine the purpose of the talk exchange, and so whether the statement cooperated with that purpose. Let's consider three cases as examples, beginning with the criminal entering a plea of "Not guilty" for a crime he knows he committed. What is the purpose of this talk exchange? If it is for the accused to *truthfully* enter a plea, then clearly the

---

<sup>6</sup> The use of assertion in a definition does not *ipso facto* place that definition in the deceptionless category. Williams (2002) and Kupfer (1982) do so, but also include the intent to deceive condition. Hence, they do not need to explicate the notion of assertion to handle the problematic cases of irony, jokes, etc. Therefore, they lack the hallmark of the deceptionless definitions.

<sup>7</sup> Observing the maxims is, for Grice, only typically how we cooperate with the common purpose of the talk exchange because speakers can also ostentatiously flout a maxim, i.e., violate a maxim in what they literally say, but in such a way that they expect their audience to recognize that the violation was intentional, so that the audience realizes what is meant is still cooperative with the common purpose.

criminal is not cooperating with that purpose. The case of the plagiarizing student seems similar. The dean’s refusal to punish plagiarism does not matter; the *official* purpose of the inquiry is to determine if the student plagiarized. So, when the student says, “I did not plagiarize,” the student said something he believes to be false, and in so doing was uncooperative with the official purpose of the talk exchange. If we find the accepted purpose of the inquiry is something different, then the student didn’t lie. Next, consider Nixon’s utterance of “I’m not a crook.” He said something he believed to be false (let’s assume), deceitful, and uncooperative with the journalists’ purpose of discovering Nixon’s involvement in the Watergate scandal. So he lied. But the deceit had nothing to do with it (though it may make the lying morally worse).

One might argue for different accounts of the purposes of these conversations. If, for instance, the purpose of the criminal’s talk exchange is merely to enter a plea (and any plea will suffice), then the criminal has cooperated and so not lied.<sup>8</sup> Such an argument speaks to the context-sensitive nature of lying, including bald-faced lying. A speaker can utter the same believed-false statement in the two different conversational contexts (and so with two different purposes to the conversations) and only lie in one of them. So whether a bald-faced lie (even without the intent to deceive) is a lie depends on the purpose of the conversation. While we can argue over specific cases, the arguments center on what the conversational purpose is (and whether the utterance cooperated with it). Thus, if someone wishes to object that the criminal has not lied based on a different purported purpose of the talk exchange, then the point has been conceded that cooperation (or the lack thereof) is central to the notion of lying. I take it to be a strength of the cooperation-based definition that it can allow for and explain the reason for different views on whether *particular* apparent bald-faced lies are lies, while it also still allows for the possibility of bald-faced lies in general.

---

<sup>8</sup> I have no strong intuitions either way on this example, and neither view threatens the cooperation-based definition. The point of the example is to demonstrate the importance of cooperation to lying. Whether the criminal was uncooperative and so lied depends on the conversation’s purpose. It is interesting though that even if the purpose only requires any plea (instead of a truthful one), then the criminal has still violated the first maxim of Quality, which enjoins all speakers not to say what they believe to be false. This points to a tension in Grice’s theory that he never fully resolved. Is observing (or intentionally flouting) his conversational maxims necessary in order to observe the Cooperative Principle? Grice often sounds as if he thinks so. Some talk exchanges, however, do not appear to require truthfulness in order to cooperate with their purposes; some even require that you not be truthful, such as telling lies as a game.

It's worth clarifying at this point that being uncooperative is not enough by itself to count as lying. Consider the following example from Grice: "At a genteel tea party, A says *Mrs. X is an old bag*. There is a moment of appalled silence, and then B says *The weather has been quite delightful this summer, hasn't it?*" (1989, pg. 35). Though Grice doesn't specify, let's initially assume that the weather has been delightful. By not observing the maxim of Relation ("Be relevant"), B is clearly being uncooperative with A, who proposed a direction for the conversation. B refused to adopt that conversational direction and rag on Mrs. X. Yet, B didn't lie, since what B said was true (and believed to be so). Now let's suppose that the weather has in fact been positively horrid this summer. B now said something B knows to be false. B is also being just as uncooperative as before. This time, however, B is lying.

So far, it might look like my cooperation-based definition is but another deceptionless definition. Like the deceptionless class, the cooperation-based definition rejects the intent to deceive condition. Nevertheless, it is structurally different from the deceptionless definitions. In deceptionless definitions, a speaker *asserts* (or proposes for common ground or warrants the truth of) a believed-false statement. In the cooperation-based definition, on the other hand, by literally *saying* a believed-false statement a speaker fails to cooperate with a talk exchange's purpose or direction. What one literally says and what one asserts are not necessarily the same thing. As we will see, this distinction is crucial for the deceptionless definitions when it comes to jokes, metaphors, irony, and fiction. For instance, according to deceptionless definitions, when a comedian (Mitch Hedberg) *says* in a comedy club "My fake plants died because I did not pretend to water them," he did not assert (or proposes for common ground or warrants the truth of) anything about fake plants.

### 3 Loose Talk

The next hazard to navigate in defining lying is loose talk. Though loose talk is not a new topic (cf., Sperber and Wilson 1985-1986), it is an issue that has gone overlooked in the literature on lying and will cause problems for existing definitions. Loose talk represents an important class of cases because we engage in loose talk all the time. Loose talk is an utterance that is not literally true, but is close enough for the purposes of the conversation. I may tell you



it's 3:30, when my watch actually shows 3:29.<sup>9</sup> When in New York and asked about his hometown, a friend might tell you she lives in Dallas, though she really lives just outside the city limits.<sup>10</sup> You may tell me that the car's gas tank is half-full, when in truth it is a bit below that (and you know that fact). In none of these cases did anyone lie.

To narrow our focus, let's take the following example of loose talk. Suppose that Stephen Hawking says, "It's ninety-three million miles from the Earth to the Sun." This is false. Not only is that figure rounded up slightly, but also, as Kepler taught us, the Earth's orbit is elliptical, not circular, so the precise distance is not fixed. And Prof. Hawking knows all this. So he's said something he knows to be false.

Has Hawking lied? Not only has he said something he knows to be false, but he said it with the intention of getting his audience to believe something false. Thus, it would seem Hawking intended to deceive his audience. So standard deception-based definitions of lying must therefore regard his utterance as a lie since he had the intent to causally contribute to his audiences acquiring a belief in something he believes to be false (which is the first of Chisholm and Freehan's (1977) types of deception).

But what of deceptionless definitions? Let's take Stokke (2013) for example. According to Stokke, a speaker *S* lies if *S* says *p* to *H*, *S* believes *p* to be false, and *S* proposes *p* to be common ground. Stokke draws upon Stalnaker's (2002) conception of common ground as a set of beliefs shared between interlocutors accepted for the purpose of the talk exchange. Hawking certainly appears to be proposing this proposition about the distance to the sun as common ground, so that he and other conversational participants can refer back to it as the conversation continues. So, according to Stokke's deceptionless definition, Hawking has lied. Alternatively, by Carson's (2006) definition, Hawking has warranted the truth of the statement about the Earth-Sun distance. He certainly seems to have asserted it, as Fallis (2009) requires. So, deceptionless definitions will all find that he has lied. But that's a ridiculous conclusion. He didn't lie; he was speaking loosely.

---

<sup>9</sup> Van der Henst *et al.* (2002) report results from three experiments showing a widespread willingness of speakers to give the time in rounded figures when asked by a stranger for the time.

<sup>10</sup> Though not focused on lying, Sperber and Wilson (1985-1986, pg. 163) use a similar example of loose talk and assert that it "is not misleading."

One might raise an objection here that could potentially save deception-based and deceptionless definitions alike when it comes to loose talk. The objection would claim that Hawking did not mean that the Earth-Sun distance is *exactly* 93,000,000 miles. The argument would be that Hawking is observing Grice's second maxim of Quantity (which requires speaker not to be more informative than the purpose demands). In this case though, this maxim conflicts with the first maxim of Quality. Were Hawking to say something he believes to be true (namely the precise Earth-Sun distance), it would be too informative for the purpose of this particular talk exchange. According to Grice (1989), when two maxims clash such that both cannot be simultaneously observed, a speaker can violate one but still be cooperative because the speaker observed the other maxim. The audience can recognize that both maxims couldn't have been fulfilled, and thereby works out the speaker's implicature.

Returning to our example then, the objection would claim that Hawking could not satisfy both maxims of Quality and Quantity at the same time, given the purpose of the talk exchange. Thus, Hawking has opted to violate the second maxim of Quality in order to satisfy the first maxim of Quantity, and so he conversationally implicates that the Earth-Sun distance is *approximately* 93,000,000 miles. Likewise when I said that the time is 3:30, I implicated that it is *approximately* 3:30. Let's call this the **approximate objection**. Loose talk would then be saying something the speaker believes to be false, but implicating something the speaker believes to be true. So, loose talk would not be a problem for either definitions of either category. The deception-based theorists could claim that because of the approximate implicature, those engaged in loose talk lack a deceptive intent. The deceptionless definitions could claim that in loose talk, one does not assert something believed to be false, due to the approximate implicature.

It seems likely that sometimes the suggested conventional implicature is present in apparent loose talk cases. However, it is far from clear that it is *always* present, and so the approximate objection fails. Consider the following counter-example. Suppose Albert Einstein one day was speaking in an introduction to physics class full of non-science majors. He tells them, "The speed of light is 300,000,000 m/s." He has engaged in loose talk. He of course knows that it is actually 299,792,458 m/s. The approximate objector would protest that Einstein was implicating that this was only an approximate speed. As a matter of fact, Einstein has

engaged in loose talk twice over. He's given the *approximate* speed of light *in a vacuum*, though he has failed to indicate explicitly either looseness in his utterance.

The problem for the approximate objection is that in cases of an implicature arising from clashing maxims, the speaker intends for the audience to work out the speaker's implicature. It would seem odd to suggest that an audience of scientific neophytes could work out that Einstein was implicating that this is the approximate speed of light. Neither does it seem plausible to suggest that Einstein would expect such an audience to work out that he meant *in a vacuum*.<sup>11</sup> These points do not *ipso facto* mean that Einstein could not have implicated that it was the approximate speed of light in a vacuum, since the presence of an implicature is determined solely by the intentions of the speaker. It strains credulity, however, to think that Einstein would have such an intention. This counterexample is not a far-fetched one. At least some cases of genuine loose talk exist without an approximate implicature to save them. Therefore, the approximate objection fails. Both deception-based and deceptionless definitions will then mistakenly find genuine, implicature-less cases of loose talk to be lies.

With the approximate objection overcome, the precise nature of loose talk is clearer. For Sperber and Wilson (1985-1986), loose talk is like metaphorical speech in that both are non-literally true utterances that lack an implicature, and loose talk is simply looser than metaphors. If correct, then loose talk belongs in same category as metaphor, irony, jokes, and fiction (which we will consider below). If we look closely, however, we can see a crucial difference between metaphor and loose talk. Metaphorical speakers intend for the looseness of their metaphors to be realized. When Burns says, "My love is a red, red rose," he expects his reader to realize that his love is not *literally* a rose. Those engaged in loose talk, on the contrary, do not expect that the looseness of their utterance to be realized. Because of this difference in speaker's intentions, loose talk is a distinct problem for any definition of lying.

When it comes to loose talk, the cooperation-based definition succeeds where other definitions fail. Hawking did not lie about the Earth-Sun distance. I did not lie when falsely

---

<sup>11</sup> The approximate objector might counter that Einstein is not specifically implicating this is the speed of light *in a vacuum*, but rather that this is the speed of light *in some unspecified condition(s)*. This line of reasoning is equally problematic. Einstein has no reason to expect that these scientific fledglings would suspect that the speed of light could vary under different conditions. Hence, he would not expect them to be able to work out even such a vague implicature as suggested. Therefore, Einstein is unlikely to have implicated it.

saying the time is 3:30. Einstein did not lie though omitting to mention anything about vacuums. The reason no one lied in these examples of loose talk is that in all of them, the speakers were nevertheless being cooperative. As the approximate objection suggested, in cases of loose talk there is a clash between two maxims, typically the first maxim of Quality (“Do not say what you believe to be false,”) and the second maxim of Quantity (“Do not make your contribution more informative than is required”) (Grice 1989, pg. 26). The mistake the objection made, however, was in claiming that this clash was the source of the alleged approximate implicature. Instead, I argue that the clash only explains how the speakers are being cooperative.

Consider the contexts in which both Einstein and Hawking are speaking. They are not giving technical talks requiring extreme precision. Rather, the purpose of both talk exchanges is to give a general overview of a topic to an audience of scientific novices. Too much specific information would confuse or intimidate the audiences, as Hawking and Einstein are aware. So they use rounded figures and did not provide all the relevant information. Neither of them intends nor expects their audiences to realize that there are two conflicting maxims at play. They do not expect their audience to work out (or even be able to work out) that they have rounded their figures or suppressed relevant, but complicating conditions. By opting to observe the second maxim of Quantity, both Hawking and Einstein cooperated with the purposes of their respective talk exchanges.

#### 4 Addressee

To whom can one lie? Most theorists contend that lies are told *to someone*, namely the person(s) to whom one is addressing. This is an intuition shared by most subsequent definitions of lying. Simpson (1992) provides a good example here. Suppose that Alice is speaking to Bruce, both of whom know that Clive is eavesdropping. Conspiring against Clive, Alice says to Bruce, “There are no police on the road,” though both she and Bruce believe there are police on the road. Simpson argues that even though Alice said something she believes to be false, she has not lied to Clive or to Bruce. According to Simpson (and most deception-based definitions), since Alice wasn’t trying to deceive Bruce (because she knew he also believed the police were present), she did not lie to him. Additionally, though Alice was trying to deceive Clive, she was not addressing him, and so did not lie to him either. Deceptionless definitions can arrive at the same

conclusion by a different route. Alice did not assert (or something similar) *to Clive* that the police were not on the road, since one only asserts something to the intended addressee. Neither did she assert it to Bruce, since she did not intend for Bruce to think that she intended Bruce to think that she believes there are no police on the road. Hence, according to deceptionless definitions as well, Alice did not lie.<sup>12</sup>

The cooperation-based definition similarly rejects the possibility of lying to an eavesdropper, stipulating that lies can only be made to addressees. Alice and Clive are not having a talk exchange, and so there is no common purpose with which Alice may cooperate or not. Hence, she did not lie to Clive. Neither did she lie to Bruce, though having said something she believed to be false. The purpose of Alice and Bruce's talk exchange (or at least one purpose of possibly many) was to deceive Clive. As such, saying something she believed to be false clearly cooperated with that purpose. Therefore, she did not lie to Bruce.

In general then, according to the cooperation-based definition, one only lies to someone when there is a purpose between them with which one does not cooperate. There is no such purpose with an eavesdropper. Likewise, one cannot lie to things that cannot have purposes. Though I may speak to my car, there is no mutual purpose we can share with which we can cooperate. So, if I say ostensibly to my car, "You've got enough gas to make it home," while not believing it, I have not lied to my car.

One might object at this point that by the cooperation-based definition, an author (especially a long-dead one) cannot lie to her or his readers, since when the author writes the future readers do not yet have any purpose with which the author can cooperate or not. The purposes of interlocutors, however, do not have to be temporally synchronous. While the author and readers do not have a purpose *at the same time*, an author can lie to the readers. An author initially sets a purpose or direction, which the readers later adopt. If an author writes something he or she believes to be false that does not cooperate with that purpose, then the author has lied to the readers.

## 5 Implying

---

<sup>12</sup> There are a few exceptions. Newey (1997) believes that one can lie to an eavesdropper. Shibles (1985) contends that a lie does not need any addressee.

Can a speaker insinuate or imply a lie without saying it? It should come as no surprise that sometimes speakers say something that is strictly speaking true, but still intend for their audiences to understand them to mean something false. Adler (1997) offers a widely referenced biblical example. In Genesis 20, Abraham is traveling with his wife Sarah (who is also his half-sister). He fears that if people know Sarah is his wife, they will kill him. So Abraham tells Abimelech that Sarah is his sister, implying that Sarah is *only* his sister and not his wife. The question is then whether Abraham lied. To answer that question, we first need a theory of speaker meaning that can account for how speakers imply more than they say. Like Adler (1997), we can adopt Grice's (1989) well-known theory of conversational implicature. This type of implicature relates back to Grice's Cooperative Principle and conversational maxims of Quality, Quantity, Relation, and Manner, which we encountered earlier. Consider the following conversation offered by Grice (1989, pg. 32):

- A:* I am out of petrol.  
*B:* There is a garage round the corner.

*B* has implicated that he believes petrol is (likely) to be found at the open garage round the corner. As Grice explains, *A* is able to work out *B*'s meaning by assuming that *B* is being cooperative, and so what *B* means must be relevant to the talk exchange's purpose (i.e., *B* has not infringed upon the maxim of Relation). Audiences are warranted in making the assumption that speakers are being cooperative, says Grice, not just because speakers regularly do observe the Cooperative Principle, but also because it is something reasonable for speakers to follow it. The Cooperative Principle is a dictate of rationality, according to Grice. Speakers regularly implicate something more than or different from what they literally said. And audiences are generally able to understand what was implicated by assuming speakers are being cooperative.<sup>13</sup>

We can now address the question of whether Abraham lied when he *says* that Sarah is his sister (which is true), but *implicates* that Sarah is not his wife (which he knows is false). Let's

---

<sup>13</sup> Saddock (1978) and Davis (1998) misconstrue Grice on this point, and we should be careful not to make the same mistake. They contend that a conversational implicature does not exist when it cannot precisely be worked out by the audience. For instance, it is sometimes unclear what exactly a speaker conversationally implicated, because multiple cooperative interpretations of the utterance are possible. Saddock and Davis claim that in such cases, there is no implicature. Unfortunately, Saddock and Davis have found a kind of error made by speakers and taken it to be an error in Grice's theory. On Grice's intention-based account, the presence of an implicature is solely determined by the intention of the speaker. Whether the audience can figure out what was implicated is a separate, epistemic matter.

suppose that Abimelech gestures towards Sarah and asks Abraham who she is. Abraham, afraid for his life if Abimelech knew Sarah was his wife, says, “She is my sister.” Abraham has exploited Grice’s first maxim of Quantity, which states, “Make your contribution as informative as is required (for the current purpose of the exchange)” (1989, pg. 26). The quantity of information required by the purpose of the talk exchange demanded more information from Abraham, namely that Sarah was not only his (half) sister, but also his wife. So he has implicated that Sarah is not his wife, which Abimelech can work out by assuming that Abraham is being cooperative and observing the first maxim of Quantity.

Adler (1997) contends that Abraham did not lie, since what he actually said was true, a view which accords with most commentators going back to Augustine (c395a). While Adler recognizes that Abraham was being deceptive with his false implicature, the deception was of a different sort than the deception of lying. To lie, one must literally *say* something believed to be false. Following Adler, Fallis (2009) argues that Abraham has not asserted something he believes false; he’s merely implicated it. Implicating is not the same as asserting, and to lie one must assert something believed false.

There is no doubt that Abraham was being uncooperative. So, it might seem that according to my cooperation-based definition, Abraham was lying, contrary to the conclusion of Adler, Augustine, and most other commentators. That would be incorrect. I agree that Abraham did not lie. A crucial feature of the cooperation-based definition isn’t just that a speaker is being uncooperative, but by what means the speaker is being uncooperative. Recall that (C3) requires that a liar be uncooperative by having *said* something believed to be false. Abraham is being uncooperative by *implicating* something believed false, and so condition (C3) is not met. Condition (C2) is also not met, since Abraham literally said something he believes to be true. So, by the cooperation-based definition, Abraham has not lied, and neither do others who similarly say something true but implicate something believed to be false.

## 6 Jokes, Irony, Metaphor, and Fiction

The fifth definitional hazard we must navigate covers a wide range of cases, including instances of jokes, irony, metaphor, hyperbole, fiction, and play-acting. While seemingly disparate, all the relevant cases include an important common feature. They all involve a speaker saying something

believed to be false, but without thereby being uncooperative. An acceptable definition of lying needs to explain why these cases aren't lies.

Consider the following example, previously given by Fallis (2009). Han Solo says to Princess Leia, "The garbage chute was a really wonderful idea. What an incredible smell you've discovered!" Anyone familiar with *Star Wars* knows that Solo believes going down the garbage chute was a terrible idea and the only thing incredible about the smell is how bad it is. Another example is Groucho Marx's famous line "One morning I shot an elephant in my pajamas. How he got in my pajamas, I don't know."<sup>14</sup> Marx, of course, was only joking; there was no elephant.

Are Solo and Marx lying? They have said something they believe to be false. Yet, it seems highly dubious to claim that either has lied. Deception-based definitions have a simple way of explaining why not. Neither Solo nor Marx intended to deceive their audiences with their false statements. Solo actually meant that the garbage chute was a terrible idea. He relies upon his audience to realize that he said something believed to be false and so implicated something different, and he expects his audiences to work out his intended meanings. The deception-based theorists point out that this difference in intent is critical. Unlike liars, ironic speakers intend the falsehood to be recognized, and so do not intend to deceive. (Jokers like Marx might intend to deceive their audiences, but only momentarily, until the initial deception is revealed for humorous effect.) Metaphorical speakers rely on their audiences to understand that they are talking about mere similarity, not identity. Fiction writers expect their audiences to realize that they're reading fiction.<sup>15</sup> Actors expect the audience to realize that they're watching actors play-acting.

Deceptionless definitions rely upon the context or intention to deny that ironic speakers, jokers, lovers of metaphor, etc. are lying. Stokke's (2013), for instance, introduces a distinction between official and unofficial common ground. Official common ground is common ground on the standard Stalnaker account, i.e., information accepted in common by interlocutors for the

---

<sup>14</sup> This statement serves as an example of a joke and play-acting. Groucho Marx is playing the character Capt. Spalding in the film *Animal Crackers*. Capt. Spalding doesn't believe the first part of the statement "One morning I shot an elephant in my pajamas," to be true. Neither does Groucho Marx as himself. Yet neither Marx nor Capt. Spalding lied.

<sup>15</sup> It could reasonably be argued that if a fiction writer does not intend or expect the audience to realize that they're reading fiction (such as was the case with the yellow journalists of the Spanish-American War) then he does in fact lie. Such cases are also not likely fiction anymore.



purpose of a talk exchange. Unofficial common ground is similar information, but only temporarily kept as common ground and then dropped, often before the talk exchange is over. For instance, the set up of a joke is unofficial common ground, since after the punch line, neither speaker nor audience keeps the set-up as part of their common ground. Stokke contends that for cases of joking, story telling, play-acting, etc., the false assertion is only unofficial common ground, and therefore not a lie.<sup>16</sup>

The cooperation-based definition of lying handles cases of ironic statements, jokes, metaphors, fiction, etc. by noting that though what was said is believed to be false, the speaker is nevertheless being cooperative when making these utterances. Grice himself accounts for cases of irony. As he explains, an ironic speaker clearly flouts the first maxim of Quality by saying something she believes to be false (1984, pg. 34). Han Solo expects Luke and Leia to recognize that he has violated this maxim. Since they can assume he is still being cooperative, they can work out that his utterance was meant ironically. The ironic meaning was cooperative with the talk exchange's purpose. Hence, he did not lie. Similarly in cases of cases of metaphors, hyperbole, and simile, though (C1) is satisfied by the speakers saying something believed false, (C3) is not satisfied since the speakers are not *thereby* uncooperative. The presence of conversational implicatures means that these speakers are cooperating. When metaphorical speakers say something like "All the world's a stage" they've said something literally false. But on the standard Gricean account (Grice 1989, Neale 1992), they've implicated something believed to be true, and are thereby cooperative.

Alternatively, case of jokes, fiction, play-acting, etc. don't include a believed-true implicature. Nevertheless, jokers and fiction writers are being cooperative with the purpose of their talk-exchanges. Marx's utterance was made in a context in which a joke, even if unexpected, would be found agreeable by his audience. So Marx cooperated with that purpose by saying something false. Fiction necessarily involves saying something believed false. Dostoevsky cooperates with that purpose when writing the imagined tale of three fictional brothers. When someone tells the tall tale of the cowboy Pecos Bill, the storyteller intends for the audience to realize that the character is a legendary, larger-than-life figure of western lore. This intention

---

<sup>16</sup> While this distinction between official and unofficial common ground theoretically sufficient to handle these cases, my primary worry here is that it looks *ad hoc*.

means that the speaker is being cooperative, and so is not lying. All these cases are similar in that they aren't lies, despite the speakers saying something believed false, because the speakers were still being cooperative.

## 7 Cooperation and Confusion

We have now navigated our way through all five definitional hazards. Before reviewing the cooperation-based definition's success versus that of the deception-based and deceptionless definitions, it is worth elaborating further on what it means for a speaker to be uncooperative in a talk exchange. Earlier, I argued that speakers engaged in loose talk are still being cooperative. Now let's imagine alternate contexts for Hawking's statement. Suppose I ask him for the Earth-Sun distance and I need it exactly, but I failed to make clear the need for such precision. He justifiably (but falsely) believes the purpose of our talk exchange to require less precision and so gives his rounded response as before. Was Hawking uncooperative? To answer this question, we need to further explicate the nature of cooperation, since the answer will determine whether he lied according to the cooperation-based definition.

As it turned out, Hawking's rounded off answer was not as informative as the purpose of the talk exchange required. So, it might seem that he was therefore uncooperative, and hence lied. This conclusion, however, is counter-intuitive, and with good reason. It would mean that speakers lie all the time without knowing it when they misunderstand the purpose of their talk exchanges (and say something they believe false to comply with what they believed to be the purpose). Yet that conclusion cannot be right. It is nonsensical to suggest that speakers lie in cases like this one without even knowing it. Therefore, it appears that whether or not a speaker is cooperative cannot be determined solely by some fact external to the speaker. Instead, a speaker is cooperative if and only if the speaker intends his or her utterance to contribute to what the speaker believes to be the purpose of the talk exchange. Hawking did intend to contribute to what he mistakenly believed to be the purpose of the talk exchange. We cannot fault him for his misunderstanding. Therefore, he was cooperative and did not lie.

As it turns out, this emphasis on intention requires that we slightly modify the cooperation-based definition, specifically condition (C3), which states, "By saying the believed-false  $p$ ,  $S$  is being uncooperative with the purpose or direction of the talk exchange in which  $S$  is

engaged with *H*.” Phrased in this way, a speaker is being uncooperative by saying something believed to be false, not based on the speaker’s intention to be uncooperative. Therefore, I propose the following minor change:

Speaker *S* lies to audience *H* if and only if:

- (C1) *S* says *p* to *H*,
- (C2) *S* believes *p* is false,
- (C3\*) *S* intends to be uncooperative with the purpose or direction of the talk exchange in which *S* is engaged with *H*, and so says *p*.

By placing the emphasis on the intent to be uncooperative, it is now easy to see how the cooperation-based definition can handle the case of the confused politician, as raised by Carson (2006). The story goes that a politician was to give a humorous speech one day at a festive banquet and a serious policy speech another day at a formal banquet. She confuses the dates for the events and gives the policy speech at the festive banquet and the humorous speech at the formal banquet. At the formal banquet, she tells a story of the President having “broken wind” in a meeting with other heads of state. She believes this claim to be false and meant it as a joke. The audience, expecting a serious speech, did not interpret the utterance as a joke. Thus, her utterance actually does not cooperate with the purpose of a serious policy speech. Yet, the politician intended to be cooperative with that what she took to be the talk exchange’s purpose. So, I agree with Carson that in this case the politician did not lie, despite knowingly having made a false claim about the President’s flatulent faux pas. The cooperation-based definition reaches this conclusion by the lack of an uncooperative intent. We can fault the politician for miscommunicating, but not for lying.

Consideration of joking in fact reveals the uncooperative nature of lying. When accused of lying, it’s not uncommon for the accused to reply, “But I was only joking.” The accused doesn’t deny having said something false. What she is appealing to is her intention to have been cooperative. Let’s assume that the speaker really was joking, and not just trying to cover the lie. She expected the joke to go over better; she thought her audience would *retroactively* agree with her humorous purpose, though they in fact didn’t. This *post hoc* lack of consent to the speaker’s purpose doesn’t, however, render her utterance uncooperative. It’s the intention that counts. She

intended to say something funny with which she expected her audience to agree to and cooperate with.

More interesting is the other, serious speech given at festive banquet to a crowd expecting jokes. This time the politician invents “a false, but humorous story to discredit a political figure before an election,” whom we’ll suppose for ease of reference is the President again (Carson 2006, pg. 296). The politician meant to lie. Her audience though took her utterance as a joke. Carson argues that though the politician intends to lie, she fails to warrant the truth of her statement due to the context, and so she did not lie. According to his view, one can only lie if one intends to warrant the truth of a believed-false statement and *actually does* warrant the truth of that statement in the context of its utterance. I disagree. She lied, because she was not cooperating with what she believed to be the purpose of the serious talk exchange when she made the false utterance. She violated the first maxim of Quality by saying something she believed to be false, but she did not intend for her audience to recognize that violation. She intended to deceive her audience and baselessly discredit the President. Therefore, there was no implicature to render her utterance cooperative. So she lied. It does not matter that her audience at the festive banquet took her comment to be a lame joke.

## 8 Conclusion, or Finding Safe Harbor

Attempting to define lying is to set sail in treacherous waters. Charting a successful course must avoid several hazards. To date, all definitions plotted can be categorized in one of two kinds of passages, deception-based and deceptionless. The iceberg of irony, metaphor, and jokes threaten a Titanic-like end, but both groups of definitions took passages far enough south to remain afloat. Deception-based definitions, however, arguably are dashed upon the cliffs of the intent to deceive, if we countenance bald-faced lies. Deceptionless definitions appeared to be in the clear. Then we discovered a new navigational hazard, the previously uncharted lee shore of loose talk. Unprepared for the danger, deceptionless definitions crashed and sank. Even if bald-faced lies aren’t the hazard they appear, deception-based definitions still would run aground here. The cooperation-based definition alone has proven able to navigate all the definitional hazards and find safe harbor.

There is one additional uncharted concern. Like most previous attempts at defining lying, I have avoided discussion of the hidden reefs of the morality of lying. Obviously, the topic of the moral permissibility of lying is as vast and intricate as the Great Barrier Reef, so we cannot canvas all of it here. It is a fair question whether the definition of lying must also capture the reason for its moral impermissibility. The deception-based definitions offer some assistance here. If we assume that *ceteris paribus* deception is morally impermissible and lying necessarily involves the intent to deceive, then we have some idea why lying in general is morally impermissible. Yet, bald-faced lies seem equally wrong, though they include no such intent to deceive. Deceptionless definitions are capable of explaining why bald-faced lies are lies, but not why they are impermissible. In fact, definitions of this type make it harder to identify where the moral wrongness of lying of any kind resides. I take it that these objections are not singularly sufficient to disprove deceptionless definitions of lying, though they do give reason to pause, especially if an alternative theory is available. The cooperation-based definition provides such an alternative. *Prima facie*, people would seem to be morally obligated to cooperate with others in certain circumstances, or at least not be uncooperative. Lying is a particular way in which people can be uncooperative. Uncooperative speech in general is morally impermissible, though there may be times when being uncooperative as a speaker is morally justified. Therefore, my cooperation-based definition of lying suggests that a fuller exploration of the morality of cooperation is called for, specifically under what conditions a speaker is obligated to cooperate. Such a moral analysis of cooperation then has the potential to radically reinterpret Grice's Cooperative Principle and theory of conversational implicature.

## References

- Adler, J. (1997). Lying, deceiving, or falsely implicating. *Journal of Philosophy*, 94, 435–452.
- Augustine. (c395b). *Lying*. In R. Deferrari (Ed.), *Treatise on Various Subjects* (Vol. 16, 53–120). New York: Fathers of the Church.
- Bok, S. (1978). *Lying: Moral Choice in Public and Private Life*. New York: Random House.
- Carson, T. L., R. E. Wokutch, & Murrmann, K. F. (1982). Bluffing in Labor Negotiations: Legal and Ethical Issues. *Journal of Business Ethics*, 1: 13–22.
- Carson, T. L. (1988). On the Definition of Lying: A reply to Jones and revisions. *Journal of Business Ethics*, 7, 509–514.
- Carson, T. L. (2006). The Definition of Lying. *Noûs* 40, 284–306.
- Chisholm, R. M., & Feehan, T. D. (1977). The intent to deceive. *Journal of Philosophy*, 74, 143–159.
- Coleman, L. & Kay, P. (1981). Prototype Semantics: The English Word *Lie*. *Language*, 57, 26–44.
- Davidson, D. (1980). ‘Deception and Division’, in J. Elster (ed.) *The Multiple Self*. Cambridge: Cambridge University Press. 79–92.
- Davis, W. A. (1998). *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*. Cambridge: Cambridge University Press.
- Fallis, D. (2009). What is Lying? *Journal of Philosophy*, 106 (1), 29–56.
- Fallis, D. (2010). Lying and Deception. *Philosophers’ Imprint*, 10 (11), 1–22.
- Frankfurt, H. G. (1988). On Bullshit. In *The Importance of What We Care About*. Cambridge: Cambridge University Press. 117–133.
- Grice, P. (1957). Meaning. *Philosophical Review*, 66, 377–88.
- Grice, P. (1969). Utterer's meaning and intentions. *Philosophical Review*, 78, 147–77.
- Grice, H. P. (1989). *Studies in the Ways of Words*. Cambridge, Mass.: Harvard University Press.
- Jones, G. (1986). Lying and intentions. *Journal of Business Ethics*, 5, 347–349.
- Kemp, K. W. & Sullivan, T. (1993). Speaking Falsely and Telling Lies. In *Proceedings of the American Catholic Philosophical Association* 67: 151–170.
- Kupfer, J. (1982). The Moral Presumption Against Lying. *Review of Metaphysics*, 36, 103–126.
- Neale, S. (1992). Paul Grice and the Philosophy of Language. *Linguistics & Philosophy*, 15, 509–559.

- Newey, G. (1997). Political Lying: A Defense. *Public Affairs Quarterly*, 11, 93-116.
- Sadock, J. (1978). On Testing for Conversational Implicature. In P. Cole (ed.), *Syntax and Semantics, Vol. 9: Pragmatics* (281–298). New York: Academic Press.
- Schiffer, S. (1972). *Meaning*. Oxford: Clarendon Press.
- Shibles, W. (1985). *Lying: A Critical Analysis*. Whitewater, Wisconsin: The Language Press.
- Simpson, D. (1992). Lying, Liars and Language. *Philosophy and Phenomenological Research*, 52, 623-639.
- Sorensen, R. (2007). Bald-Faced Lies! Lying Without The Intent To Deceive. *Pacific Philosophical Quarterly*, 88, 251-264.
- Sperber, D. & Wilson, D. (1985-1986) Loose talk. *Proceedings of the Aristotelian Society* LXXXVI, 153-71.
- Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25, 701-721.
- Stokke, A. (2013). Lying and asserting. *Journal of Philosophy*, 110 (1), 33-60.
- Van der Henst, J.-B., Carles, L., & Sperber, D. (2002). Truthfulness and Relevance in Telling Time. *Mind and Language*, 17 (5), 457-466.
- Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.