

What is Lying?

It's not a lie if you believe it.
– George Costanza

A lie with a nod is still a lie, but it's an easy lie.
– Philip Marlowe

1. Introduction

The main question that we typically have about lying is whether someone (a politician, a witness, a spouse, etc.) is lying to us. And we frequently expend a lot of energy to find out. For example, the police use polygraphs that monitor heart rate, perspiration, and other physiological indicators of lying on criminal suspects. And, recently, researchers have even used fMRI to find that different areas of the brain are activated depending on whether someone is lying or telling the truth.¹

But there are also many important philosophical questions about lying. For example, many moral philosophers have studied the morality of lying.² Most notably, Immanuel Kant (1959 [1785]) argued that it is always wrong to lie by asking us to

¹ See, e.g., Feroze B. Mohamed, Scott H. Faro, Nathan J. Gordon, Steven M. Platek, Harris Ahmad, and J. M. Williams, "Brain Mapping of Deception and Truth Telling About an Ecologically Valid Situation: Functional MR Imaging and Polygraph Investigation—Initial Experience," *Radiology*, CCXXXVIII (2006):679-88. There are also techniques that can be used to identify lies in recorded information. For example, researchers have used textual analysis to find that liars are somewhat less likely to use first-person pronouns. See, for example, Matthew L. Newman, James W. Pennebaker, Diane S. Berry, and Jane M. Richards, "Lying Words: Predicting Deception From Linguistic Styles," *Personality and Social Psychology Bulletin*, XXIX (2003):665-75.

² For example, Plato, *Republic*, C. D. C. Reeve, trans. (Indianapolis: Hackett, 2004 [360 BCE]), Augustine, "Lying," in *Treatises on Various Subjects*, vol. XVI, Roy J. Deferrari, ed. (New York: Fathers of the Church, 1952 [395]), pp. 53-120, Thomas Aquinas, "Of Lying," in *Summa Theologica*, vol. XII, (London: Burns, Oates, & Washbourne, 1922), pp. 85-98, Immanuel Kant, *Foundations of the Metaphysics of Morals*, Lewis W. Beck, trans. (New York: Macmillan, 1959 [1785]), Sissela Bok, *Lying*, (New York: Random House, 1978), Jonathan E. Adler, "Lying, Deceiving, or Falsely Implicating," this *Journal*, XCIV (1997):435-52. In fact, this is arguably a central issue in ethics. For example, according to Kant, lying is "the greatest violation of a human being's duty to himself regarded merely as a moral being" (quoted in Bernard Williams, *Truth and Truthfulness*, (Princeton: Princeton, 2002), p. 102).

imagine what would happen if everybody lied when it was to their advantage. More recently, Sissela Bok (1978) argued that it is wrong to lie more often than we think because we often underestimate the personal and social costs of lying.

In addition, many epistemologists are interested in lying.³ A great number of our beliefs about the world are based on what other people tell us. But it is not clear whether these beliefs count as knowledge if these other people might be lying to us.

But in order to answer such questions about lying in ethics and epistemology, we first need to know what it means for something to be a lie.⁴ In other words, philosophers need to define their terms. Saying what it means for something to be a lie is the goal of this paper.

2. Conceptual Analysis

In order to provide such a definition, I will be engaging in *conceptual analysis*.⁵ Plato famously used this technique in his dialogues to try to understand such concepts as justice, knowledge, and love. And it has been an important part of philosophical practice ever since. (Harry Frankfurt's attempt to understand the concept of *bullshit*—by

³ For example, Ludwig Wittgenstein, "Notes for Lectures on "Private Experience" and "Sense Data", " *Philosophical Review*, LXXVII (1968 [1935]):275-320, here p. 280, Tyler Burge, "Content Preservation," *Philosophical Review*, CII (1993):457-88, here p. 474, Peter J. Graham, "What Is Testimony?," *Philosophical Quarterly*, XLVII (1997):227-32, Jennifer Lackey, "Norms of Assertion," *Nous*, XLI (2007):594-626, here p. 602, Paul Faulkner, "On Dreaming and Being Lied to," *Episteme*, II (2006):149-59, Dan O'Brien, "Testimony and Lies," *Philosophical Quarterly*, LVII (2007):225-38. In fact, just like Kant, Faulkner (2006) recently asked us to imagine what would happen if everybody lied when it was to their advantage. Basically, whereas Descartes was essentially worried about possibly being a character in the *Matrix*, Faulkner is worried about possibly being a character in the *Truman Show*.

⁴ Cf. Thomas L. Carson, "The Definition of Lying," *Nous*, XL (2006):284-306, here p. 284.

⁵ See, e.g., Frank Jackson, *From Metaphysics to Ethics*, (New York: Oxford, 1998), Eric Margolis and Stephen Laurence, "Concepts," *Stanford Encyclopedia of Philosophy*, (2006): <http://plato.stanford.edu/entries/concepts/>.

contrasting it with lying—is a notable, contemporary example of conceptual analysis.⁶) However, over the last several years, the legitimacy of conceptual analysis has been called into question.⁷ Thus, it will be useful to clarify exactly what the goals and methodology of this paper will be.

For each concept that he studied, Plato tried to develop a *definition* that correctly classified things as falling under that particular concept or not. Ideally, he tried to identify *necessary* and *sufficient* conditions for whether things fall under that concept (cf. Margolis and Laurence 2006, section 2.1). In our case, we would like a definition that rules in everything that is a lie and that rules out everything that is not a lie.⁸

In order to determine whether a proposed definition is correct, however, we need to have fairly reliable intuitions about whether particular (often hypothetical) cases fall under the given concept. Epistemologists, for example, typically test proposed definitions of knowledge in just this way.⁹ In our case, we have to be able to tell (given that we know certain things about the beliefs and intentions of the speaker) whether specific statements are lies.

Conceptual analysis has been criticized because of this reliance on intuition. For example, several people have argued that our intuitions about hypothetical cases cannot

⁶ *On Bullshit*, (Princeton: Princeton, 2005). Given their close connection, it may be possible to give a definition of bullshit along lines very similar to the definition of lying that I propose in section 3 below (e.g., you are *bullshitting* if you say something for which you (believe you) lack adequate evidence and you believe that you are in a situation where the following norm of conversation is in effect: “*Do not say that for which you lack adequate evidence.*”). However, I will not pursue that possibility further in this paper.

⁷ See, e.g., Stephen Stich and Jonathan M. Weinberg, "Jackson's Empirical Assumptions," *Philosophy and Phenomenological Research*, LXII (2001):637-43, Jonathan M. Weinberg, Shaun Nichols, and Stephen Stich, "Normativity and Epistemic Intuitions," *Philosophical Topics* XXIX (2001):429-60, Margolis and Laurence 2006, section 5.2, Andrew Melnyk, "Conceptual and Linguistic Analysis: A Two-Step Program," *Nous*, XLII (2008):267-91.

⁸ The question of what 'lying' means is, of course, different from the question of whether someone is lying to us on a particular occasion. Even if we know what the definition of lying is, we may not know whether the definition is satisfied on any particular occasion (cf. Williams 2002, p. 97).

⁹ See, e.g., Edmund Gettier, "Is Justified True Belief Knowledge?," *Analysis*, XXIII (1963):121-23.

give us “a priori knowledge of necessary truths” (Melnyk 2008, p. 267). Also, there is empirical evidence that such intuitions are fallible.¹⁰

But even if they can only provide us with fallible, a posteriori knowledge, there are good reasons to think that our intuitions can be used effectively to test proposed definitions of important concepts (cf. Melnyk 2008, p. 288). For instance, our intuitions have to be a pretty good guide to how words are commonly used (cf. Jackson 1998). If they were not, we would have a lot of trouble communicating with each other.

Furthermore, trying to capture common usage is arguably a good way to identify useful concepts.¹¹ It would be very surprising if humans had developed terms like ‘knowledge’ and ‘lying,’ but these terms were not getting at important phenomena in the real world.

Admittedly, people from different cultures may use the same words very differently (cf. Weinberg et al. 2001). But the goal in this paper is to determine how the word ‘lying’ is used in *this* culture. Thus, relying on the intuitions of members of this culture (in this case, philosophers) seems perfectly appropriate.¹²

Even if we have reliable, shared intuitions, however, we might still worry that the goal of conceptual analysis is unattainable. That is, there may not be necessary and sufficient conditions for whether things fall under a given concept. Instead there may

¹⁰ For instance, our intuitions about hypothetical cases can be influenced by the order in which the cases are presented to us. See, e.g., Stacey Swain, Joshua Alexander, and Jonathan M. Weinberg, "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp," *Philosophy and Phenomenological Research*, LXXVI (2008):138-55. Also, our intuitions can be influenced by our moral beliefs. See, e.g., Joshua Knobe, "The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology," *Philosophical Studies*, CXXX (2006):203-31.

¹¹ Cf. John L. Austin, "A Plea for Excuses," *Proceedings of the Aristotelian Society*, LVII (1956):1-30, here p. 8. This is not to say that conceptual analysis is the only way to identify useful concepts. Nor is it to say that conceptual analysis is guaranteed to identify useful concepts.

¹² Even within the same culture, different people may have conflicting intuitions about particular cases (cf. Carson 2006, p. 301). This might indicate that there is more than one concept in play. However, while there are certainly some conflicting intuitions about lying (cf. footnote 17 below), there seems to be fairly broad agreement on the central cases. Even so, it would be useful to perform an empirical study (a la

simply be prototypical instances of the concept with different things falling closer to or further from these prototypes.¹³ However, capturing common usage is only one desideratum for a definition of lying. Since we are looking for a definition for philosophical use, we are also concerned with the simplicity and theoretical utility of the definition. Also, while it may not be possible to capture common usage perfectly with a concise definition of lying, some proposed definitions will do better than others. In this paper, I will offer a definition of lying that captures common usage better than the other proposed definitions.¹⁴ I will argue that the other proposed definitions are either too broad (that is, they rule in some things that are clearly not lies) or too narrow (that is, they rule out some things that clearly are lies).

3. My Definition of Lying

There is an informal tradition in the conceptual analysis literature of sneaking up on the official definition of some concept. That is, one starts with an obvious candidate definition, says why it won't quite work, tweaks the definition, says why it still won't work, tweaks the definition again, etc. By the time the reader reaches the official definition, she is often dizzy enough to capitulate. I am going to break with tradition and simply start out by stating my official definition.

Weinberg et al. 2001 or Swain et al. 2008) to test the proposed definitions against intuitions of people who do not have a stake in this philosophical debate.

¹³ See, e.g., Stich and Weinberg 2001, pp. 638-40, Marjorie Taylor, Gretchen L. Lussier, and Bayta L. Maring, "The Distinction Between Lying and Pretending," *Journal of Cognition and Development*, IV (2003):299-323, Margolis and Laurence 2006, section 2.2.

¹⁴ It turns out that most philosophers who discuss lying are actually concerned with something more restrictive than 'lying' as the term is commonly used. In the final section of this paper, I will argue that a slight modification to my definition of lying can capture this narrower concept.

I think that you lie when you *assert* something that you believe to be false.¹⁵

Basically, you lie when you “go on the record” with something that you believe to be false (cf. Carson 2006, p. 290, Sorensen 2007, p 252).

You *lie* to X if and only if: **(AL)**

1. You assert that *p* to X.
2. You believe that *p* is false.

Of course, in order to really define lying, I should also say exactly what it means to assert something. 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* (viz., “Do not say what you believe to be false”) is in effect as a *norm of conversation*.¹⁶

Plugging this account of assertion into **AL**, we get my definition of lying.¹⁷

You *lie* to X if and only if: **(BNL)**

1. You state that *p* to X.

¹⁵ Cf. Roderick M. Chisholm and Thomas D. Feehan, "The Intent to Deceive," this *Journal*, LXXIV (1977):143-59, here p. 152, Adler 1997, p. 435, Gary Watson, "Asserting and Promising," *Philosophical Studies*, CXVII (2004):57-77, here p. 72, Roy Sorensen, "Bald-Faced Lies! Lying Without the Intent to Deceive," *Pacific Philosophical Quarterly*, LXXXVIII (2007):251-64, here p. 256.

¹⁶ *Studies in the Way of Words* (Cambridge: Harvard, 1989), p 27. You can deceive people by violating several different norms of conversation (cf. Grice 1989, p. 30, D. S. Mannison, "Lying and Lies," *Australasian Journal of Philosophy*, XLVII (1969):132-44, here p. 132). But you need to violate Grice’s first maxim of quality, in particular, in order to be lying. I am assuming that, if this norm is in effect and you do not act in accordance with it, then you have violated it. Thus, a simpler statement of my definition would be that you are lying if you believe that you are violating Grice’s first maxim of quality.

¹⁷ In this paper, I argue that my definition of lying gets all of the central cases right. But there are many borderline cases where intuitions are less certain (cf. Carson 2006, p. 294, Sorensen 2007, pp. 258-59). In several footnotes below, I describe how my definition can easily be modified—in a way that preserves the essential insight that lying has to do with violating a particular norm of conversation—to accommodate conflicting intuitions on such cases. For example, if Graham Priest states that “This sentence is false” in a situation where he believes that Grice’s first maxim of quality is in effect, it is not clear that he is lying. Priest believes that this statement is false *and* that it is true. See his "The Logic of Paradox," *Journal of Philosophical Logic*, VIII (1979):219-41. Thus, he is obeying the maxim “Try to make your contribution

2. You believe that you make this statement in a context where the following norm of conversation is in effect:

Do not make statements that you believe to be false.

3. You believe that p is false.

This definition of lying correctly rules in many prototypical instances of lying. For example, suppose that I am at a fancy party in Washington, D.C., and a beautiful woman comes up and asks who I am. I might say with complete seriousness, “I am the Prince of Denmark” in order to impress her.¹⁸ My definition gives the intuitively correct result that this is a lie. In this case, I make a statement, I believe that it is false, and I believe that this is a situation where I should not say what I believe to be false. Similarly, Silvio is lying if he testifies in court that “Tony was home with me at the time of the murder” in order to deceive the jury and get his boss acquitted.¹⁹

In addition, my definition correctly rules out many false statements that are not lies. For example, suppose that I am on stage at my local community theater and a beautiful woman comes up and asks who I am. I might then say that “I am the Prince of Denmark” because it is a line from the play that we are performing. My definition gives the intuitively correct result that this is not a lie. In this case, I make a statement, I believe that it is false, but I do not believe that this is a situation where I should not say

one that is true” (Grice 1989, p. 27). In order to accommodate the intuition that Priest is not lying, my definition might be modified to use the norm: *Do not make statements that you believe to be only false.*

¹⁸ I looked him up on Wikipedia. He is about the same age and he looks a little bit like me. So, I might get away with it. But I am lying regardless of whether I expect to get away with it.

¹⁹ Grice’s maxims are usually unstated. But when Silvio takes the witness stand and swears to tell the truth, the whole truth, and nothing but the truth, he makes it fairly explicit that he believes that Grice’s first maxim of quality is in effect. When he then lies under oath, it counts as *perjury* (cf. Stuart P. Green, *Lying, Cheating, and Stealing* (Oxford: Oxford, 2006), pp. 133-47).

what I believe to be false.²⁰ Similarly, 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 (cf. Grice 1989, p. 30, Watson 2004, p. 65).²¹

4. Norms of Conversation

Before I discuss why the other definitions of lying are wrong, there are a few important things that should be noted about the assertion condition in my definition. First, my analysis of assertion in terms of Grice’s first maxim of quality is not intended to be a *complete* analysis of assertion. I am only trying to capture the normative component of assertion that is necessary for lying.²²

Second, the assertion condition is not that you believe that you are in a situation where, *all things considered*, you should not say things that you believe to be false. For example, suppose that a homicidal maniac shows up in the audience of an

²⁰ It might be suggested that an actor on stage is not making *statements* at all. But for purposes of this paper, I do not draw any distinction between *stating* something and just *saying* something (cf. Chisholm and Feehan 1977, p. 151). It should be noted that actors *can* lie on stage (e.g., when they speak *sotto voce* to each other). But utterances that are part of the performance are never lies.

²¹ Even if a speaker sends such a signal, the audience may not pick up on it (cf. Carson 2006, p. 294). For example, an adult might tell a tall tale (i.e., an exaggerated story) to a group of very young children that he expects to take him totally seriously. Although it is clear to the speaker that Grice’s first maxim of quality is not in effect in this situation (in fact, telling children such tales ultimately helps them to learn when this norm is not in effect), some (e.g., Carson 2006, p. 298) might want to say that he is lying. In order to accommodate that intuition, my definition might be modified to simply require that a liar believe that *his audience believes* that the norm is in effect. Also, if a visitor to a foreign country does not know for sure whether a wink turns off Grice’s first maxim of quality in this unfamiliar cultural context, some might want to say that she is still lying even if she winks after saying something that she believes to be false (especially if she intends to deceive). In order to accommodate that intuition, my definition might be modified to simply require that a liar *fail to believe* that the norm is *not* in effect.

²² More substantive accounts of assertion have been proposed (cf. Robert Brandom, “Asserting,” *Nous*, XVII (1983):637-50, Watson 2004). In fact, some philosophers (e.g., Timothy Williamson, “Knowing and Asserting,” *Philosophical Review*, CV (1996):489-523) have even suggested that the relevant norm is: *Do not make statements that you do not know to be true.*

improvisational performance piece, demands that the performance continue, but threatens to shoot any performer who says something that she believes to be false. This is certainly a situation where the performers believe that they should not say things that they believe to be false. But, since this is a performance, the performers do not believe that Grice's first maxim of quality is in effect as a norm of conversation. Thus, the performers would not be lying if they did disobey the maniac.

Third, the assertion condition is not simply that you believe that you are in a situation where a norm is in effect that you should not say things that you believe to be false. For example, suppose that a powerful and beloved monarch does not like to hear false statements even if they are said in jest. As a result, it is considered impolite to say things that you believe to be false in her presence. So, if I say to her that "I am the Prince of Denmark" with a wink, I am violating a norm with the same content as Grice's first maxim of quality. But I am not lying. Because of the wink, Grice's first maxim of quality is not in effect *as a norm of conversation* (and I do not believe that it is in effect) with respect to this particular statement.²³

The notion of a *norm of conversation* is essentially taken directly from Paul Grice (1989, pp. 22-40). In addition to not saying what you believe to be false, you should "not say that for which you lack adequate evidence," you should "make your contribution as informative as is required," you should "avoid ambiguity," etc. The function of these norms (unlike norms of politeness, for example) is to facilitate communication in normal conversations.²⁴ As Grice (1989, p. 28) puts it, "all sorts of other maxims (aesthetic,

²³ The norm would typically be back in effect for any subsequent statements.

²⁴ Social norms typically provide a solution to some interaction problem (cf. Edna Ullmann-Margolit, *The Emergence of Norms* (Oxford: Oxford, 1977), Philip Pettit, "Virtus Normativa: Rational Choice Perspectives," *Ethics*, C (1990):725-55). For example, the function of norms like "Drive on the right side

social, or moral in character), ... are also normally observed by participants in talk exchanges ... The conversational maxims, however, ... are specially connected ... with the particular purposes that talk ... is adapted to serve and is primarily employed to serve.” In such contexts, we (typically) obey such norms, expect other people to obey them, and believe that other people expect us to obey them. We quickly learn when such norms are in effect (e.g., in a normal conversation) and when they are not (e.g., in the performance of a play). And we have such knowledge even if we have not read Grice and could not put our knowledge explicitly into words.²⁵ Also, when we believe that such a norm is in effect, we think that it is reasonable to rebuke someone for violating it.²⁶ But we recognize that such violations are sometimes excusable (viz., when there are other interests or norms that trump the norms of conversation).²⁷

Finally, I am not taking any position on the *moral* force of Grice’s first maxim of quality. It could be that you have a *prima facie* moral duty to follow this norm of conversation (as Kant would probably say) or it could have no more force than a rule of etiquette. In fact, it could have no moral force at all. This agnosticism has the advantage of leaving open the question of whether lying is always wrong (cf. Carson 2006, p. 288).

5. Points of General Agreement

of the road” is to keep us from running into each other. Norms of conversation, in particular, help us with the coordination problem that we face when we try to communicate with each other. And these norms apply in the case of written as well as oral communication.

²⁵ Thus, it is clearly possible for young children, as well as adults, to lie on my definition.

²⁶ In this regard, my account of assertion is very similar to accounts of assertion in terms of *guaranteeing* or *warranting* (see section 7.2 below).

²⁷ While the norms of conversation are clearly distinct from other types of norms, they are similar in many respects to the social norms that govern other everyday activities. For example, we usually know when the norm “Drive on the right side of the road” is in effect (e.g., on a city street) and when it is not (e.g., at a demolition derby). We typically rebuke people who violate the norm, but we recognize that such violations

Despite its initial plausibility, no one (as far as I know) has proposed this particular definition of lying (**BNL**). In fact, the other proposed definitions are quite different in many respects. And, as we go along, we will consider several examples that my definition gets right and that the other definitions do not. But before we get to the points of disagreement with these other definitions (and why they are wrong), it will be useful to note a few things about the points of agreement.

The first thing that everybody pretty much agrees on is that, in order to lie, you have to make a statement.²⁸ Basically, you have to use language to express a proposition.²⁹ So, for example, while a poker player who makes a very large bet may be trying to deceive you about the strength of his hand, he is not lying to you. Similarly, someone who packs his luggage to mislead you into believing that he is leaving on a trip has not lied to you.³⁰

Admittedly, the word ‘lie’ is sometimes used to refer to cases where no statement is made. For example, there are “lies of omission” where you fail to state that *p* when it

are sometimes excusable (e.g., when a person or a big rock is lying in the road). See Pettit (1990) for a discussion of the defining characteristics of social norms.

²⁸ It is important to note that the statement that you make is not always the same as the literal meaning of your words. For example, suppose that you have a PhD in Philosophy, but have no medical training. If you come across a seriously injured person on the street and say “I cannot help you; I am not a doctor,” you are not lying. And if you say “I can help you; I am a doctor,” you *are* lying. (Thanks to Jonathan Adler for this example.) The same issue arises with pretending (e.g., when my niece turns on a flashlight and threatens me with her “light saber”) and with metaphor (e.g., when I say of my messy roommate that he is a “pig”). Basically, the liar has to believe that her audience understands her statement to mean a certain thing, which she believes to be false. See Chisholm and Feehan (1977, pp. 150-51) for an analysis of exactly what it means to make a statement. But since this is a difficult issue that all of the proposed definitions of lying must address, I will set it aside for the purposes of this paper.

²⁹ You may not have to actually succeed in communicating this proposition to your audience (cf. Carson 2006, p. 299). For example, you can probably lie to someone who fails to understand what you are saying or who already believes what you have to tell him.

³⁰ Cf. Brian Huss, “Bluffing, Lying, and Bullshitting,” in Eric Bronson, ed., *Poker and Philosophy* (Chicago: Open Court, 2006), pp. 127-37, here pp. 127-30, Immanuel Kant, “Ethical Duties Towards Others: Truthfulness,” in Louis Infield, trans., *Lectures on Ethics* (New York: Century, 1930), pp. 224-35, here p. 226.

would be expected that you would have if p were true. But most philosophers agree that such deceptions are not lies, strictly speaking.³¹

It should be noted, however, that you do not have to make this statement out loud.³² For example, you can lie by writing something down (e.g., in a letter). You can lie by sending smoke signals. In fact, as Philip Marlowe points out, you can even lie just by nodding in response to a question.

The second thing that everybody pretty much agrees on is that, in order to lie, you have to believe that your statement is false. So, as George Costanza points out, if you make a statement that you believe to be true, you are not lying.³³ In fact, you are not lying even if you intend to deceive someone by making this statement.³⁴

In addition, almost all philosophers think that your statement does not actually have to be false.³⁵ Or, as Saint Augustine (1952 [395], p. 55) put it, “a person is to be

³¹ Cf. James Mahon, "The Definition of Lying and Deception," *Stanford Encyclopedia of Philosophy*, (2008): <http://plato.stanford.edu/entries/lying-definition/>, section 1.1. But contrast Paul Ekman, "Lying and Deception," in Nancy L. Stein, Peter A. Ornstein, Barbara Tversky, and Charles Brainerd, eds., *Memory for Everyday and Emotional Events* (Mahwah, New Jersey: Lawrence Erlbaum: 1997), pp. 333-47, here pp. 334-335. Such “lies” deceive by violating one of Grice’s (1989, p. 26) maxims of *quantity*. In other words, the speaker is not being “as informative as is required.” Although it is not strictly speaking a lie, if you have an obligation to state that p to X if p is true, a lie of omission can be a form of deception that is as morally and epistemically problematic as a lie.

³² Cf. Frederick A. Siegler, "Lying," *American Philosophical Quarterly*, III (1966):128-36, here p. 128, Chisholm and Feehan 1977, p. 149, Carson 2006, p. 287, Mahon 2008, section 1.1.

³³ If you innocently repeat a false statement made by someone else, we may call the statement itself a lie. But you are not lying when you pass it along (cf. Harry Frankfurt, "Reply to G. A. Cohen," in Sarah Buss and Lee Overton, eds., *Contours of Agency* (Cambridge: MIT, 2002), pp. 340-44, here p. 340).

³⁴ Cf. Augustine 1952 [395], pp. 56-59, Adler 1997, Mahon 2008, section 1.2. But contrast Barry O’Neill, "A Formal System for Understanding Lies and Deceit," (2003): <http://www.sscnet.ucla.edu/polisci/faculty/boneill/bibjer5.pdf>. Although it is not strictly speaking a lie, such a *false implicature* can be a form of deception that is as morally and epistemically problematic as a lie (cf. Adler 1997).

³⁵ See, e.g., Siegler 1966, pp. 130-32, Mannison 1969, pp. 134-37, Warren Shibles, "A Revision of the Definition of Lying as an Untruth Told With Intent to Deceive," *Argumentation*, II (1988):99-115, here p. 101, Mahon 2008, section 1.2. Kant also seems to think that lies do not actually have to be false. See his "On a Supposed Right to Lie From Altruistic Motives." in Lewis W. Beck, trans., *Critique of Practical Reason and Other Writings in Moral Philosophy*, (Chicago: University of Chicago, 1949 [1797]), pp. 346-50. Part of the reason that he thinks that it is wrong to lie to a murderer at the door is that your statement might turn out to be true. As far as I know, Carson (2006, p. 285) is the only philosopher who thinks that actual falsity is required for lying. (But he does note that this requirement can easily be dropped from his

judged as lying or not lying according to the intention of his own mind, not according to the truth or falsity of the matter itself.” In other words, what is relevant is the state of mind of the liar rather than the actual state of the world. For example, in a short story by Jean Paul Sartre, Pablo Ibbieta tried to deceive the authorities by telling them that Ramon Gris “is hidden in the cemetery.”³⁶ Pablo believed otherwise, but it turned out that Ramon was indeed hiding in the cemetery. Even though his statement was true, Pablo was still lying.

6. Intending to Deceive

It is clear that making a statement that you believe to be false is not sufficient for lying. For example, you are not lying if you follow such a statement with a wink or if it is part of a play that you are performing. But there is disagreement among philosophers about how to rule out such cases. In this section, I will discuss the standard strategy for dealing with these cases. In the next section, I will discuss a family of alternative strategies. I argue that these strategies lead to definitions of lying that (unlike my definition) get several cases wrong.

In order to rule out such cases, most philosophers require that you also intend to deceive someone with this statement.³⁷ In other words, you are lying when you make a statement that you believe to be false with the intent to deceive.

definition.) Dictionary definitions of lying often do include this requirement. But there can easily be a difference between the way that a word is commonly defined and the way that it is commonly used. In any event, it would be a simple matter to modify any proposed definition of lying, including my own, to accommodate the intuition that actual falsity is required (cf. Carson 2006, p. 285).

³⁶ *The Wall* (New York: New Directions, 1948).

³⁷ See, e.g., Augustine 1952 [395], p. 56, Mannison 1969, p. 133, Bok 1978, p. 13, Williams 2002, p. 96, Mahon 2008, section 1.4. (It should be noted that Augustine (1952 [395], p. 60) was not completely sure whether this condition was necessary for lying.) Social scientists (e.g., J. A. Barnes, *A Pack of Lies* (Cambridge: Cambridge, 1994), p. 11, Ekman 1997, p. 334) also typically include this requirement in their

You *lie* if and only if:

(IDL)

1. You state that *p*.
2. You believe that *p* is false.
3. You intend to deceive by making this statement.

As it stands, however, this standard philosophical definition of lying is too broad. That is, it rules in some things that are not lies. For example, suppose that I decide to try to convince my new acquaintance at the fancy party that I am an actor (rather than royalty). So, I take a theatrical pose and intone, "I am the Prince of Denmark." Although I have made a statement that I believe to be false with the intent to deceive her (about my occupation), I have not lied to her (e.g., about being royalty). I have only *falsely implicated* that I am an actor.³⁸ Thus, as Bernard Williams (2002, p. 96) points out, the definition needs to require that I intend to deceive with respect to the very statement that I make.

In addition, suppose that I make a false statement to a companion (who is in on my little scheme) with the goal of deceiving someone that I know is eavesdropping on our conversation.³⁹ Since I am not even speaking to this third person, it is not clear that I

definitions of lying. It is also part of most dictionary definitions of lying (cf. Carson 2006, p. 286). According to this view, you do not have to actually *succeed* in deceiving someone in order to be lying. You just have to intend to do so.

³⁸ Cf. Grice 1989, pp. 30-31, Adler 1997. Mark Twain refers to this sort of deception as a "modified lie." See his "My First Lie, and How I Got Out of It," in *The Man That Corrupted Hadleyburg* (New York: Oxford, 1996 [1899]), pp. 167-80, here pp. 173-74. Such a "lie" may be as morally objectionable as lying (cf. Twain 1996 [1899], Adler 1997). But it is not lying.

³⁹ Cf. Chisholm and Feehan 1977, p. 156. This sort of thing certainly happens in Shakespearean plays and soap operas. For example, in Act II, Scene III of "Much Ado About Nothing," several men hold a loud conversation near where they know Benedick is hiding in order to misled him into thinking that Beatrice loves him. But it also occurs in real life. For example, fake radio transmissions were sent by the Allies to deceive the Germans during World War II. See J. B. Bell, "Toward a Theory of Deception," *International Journal of Intelligence and CounterIntelligence*, XVI (2003):244-79, here p. 249.

can be lying to this person. Thus, the definition needs to require that I make the statement *to* the person that I intend to deceive.^{40,41}

You *lie* to X if and only if: (IDL*)

1. You state that *p* to X.
2. You believe that *p* is false.
3. You intend to deceive X with respect to *p* by making this statement.

IDL* does not appear to be too broad. For example, it correctly rules out statements that are believed to be false, but that are followed with a wink or that are part of a play. (In such cases, the speaker does not intend to deceive his audience.) However, this definition is too narrow. As several philosophers have pointed out, an intention to deceive is not a *necessary* condition on lying.⁴² Thus, **IDL*** incorrectly rules out many things that are lies.

Admittedly, lying very often does involve an intention to deceive. For example, Silvio clearly intended to deceive the jury when he testified that “Tony was home with me at the time of the murder.” He wants to deceive them because he wants Tony to be

⁴⁰ In a similar vein, my definition of lying rules out the eavesdropper case by requiring that you assert *p* to X. Elizabeth Fricker claims that assertions cannot be targeted at particular individuals. See her “Second-Hand Knowledge,” *Philosophy and Phenomenological Research*, LXXIII (2006):592-618, here p. 598. That is, anyone within earshot can acquire knowledge from a sincere assertion. This might seem to suggest that I am lying to the eavesdropper. But assertions do have to have an intended audience (cf. Fricker 2006, p. 596). And, in the eavesdropper case, I am not asserting anything to anyone. For example, I do not believe that Grice’s first maxim of quality is in effect when I say something to my companion (who is in on my little scheme). So, while the eavesdropper may *believe* that she has overheard an assertion, she is wrong. And, since there is no assertion, there is no lie.

⁴¹ I do not have to have any *specific* person in mind as my intended audience. For example, a journalist who makes up a news story is lying even if she does not have any idea exactly who her readers will be (cf. Mahon 2008, section 1.3). And I may not even have to make the statement to someone that actually exists. For example, a homeowner who is awakened by a suspicious noise and says “I have a rifle” may be lying even if there is no burglar (but contrast Chisholm and Feehan 1977, pp. 157-58).

⁴² See, e.g., Siegler 1966, p. 129, Shibles 1988, p. 102, Carson 2006, p. 289, Sorensen 2007. Everybody does agree that, in order to be a lie, your stating that *p* must be an *intentional action*. For example, when the Jedi Master, Obi-Wan Kenobi, caused an Imperial Stormtrooper to say, “These aren’t the droids we’re looking for,” the Stormtrooper was not lying even if he believed that this statement was false. If someone

acquitted. Similarly, I intend to deceive my new acquaintance at the party so that she will be impressed with me.

However, there are also cases of lying that do not involve an intention to deceive.⁴³ For example, suppose that it is a regular citizen (instead of a member of his crew) that has been called as a witness to testify about Tony's whereabouts at the time of the murder. Such a witness might testify falsely about Tony's whereabouts, not because he wants this gangster to be acquitted, but because he fears for his safety if he testifies truthfully.⁴⁴ In that case, while he makes a false statement, he does not intend to deceive the jury by doing so. In fact, he may actually hope that the jury will not be deceived because he will be safer if Tony is convicted and put away. Nevertheless, he is still lying to them.⁴⁵

While the witness does not intend to deceive the jury, he does at least *foresee* that they are likely to be deceived. This suggests a weaker condition on lying.⁴⁶ This condition would also rule out statements that are believed to be false, but that are followed with a wink or that are part of a play. In those cases, the speaker does not expect that his audience might be deceived.

You *lie* to X if and only if:

(FDL)

1. You state that *p* to X.

has taken control of your vocal cords, you are not lying no matter what other conditions might hold. But intending to state that *p* is not the same as *intending to deceive* by stating *p*.

⁴³ The following example is inspired by Carson (2006, pp. 289-90).

⁴⁴ Silvio might also fear for his safety if he testifies truthfully. But we can assume that he wants Tony to be acquitted as well.

⁴⁵ It might be suggested that this example is somewhat unrealistic. That is, even a regular citizen will probably intend to deceive the jury because succeeding in deceiving them is the only way that he can be sure to escape the wrath of Tony (and his crew). But I will give a more definitive counter-example in just a moment.

2. You believe that p is false.
3. You believe that, by making this statement, X may be deceived with respect to p .

But not only does a liar not have to have an intention to deceive, a liar does not even have to foresee that someone might be deceived. We can modify the witness example to show this. Suppose that the witness knows that the jury has already seen a videotape of Tony committing the murder.⁴⁷ In that case, the witness will have no expectation that the jury will believe his statement about Tony's whereabouts at the time of the murder. Nevertheless, he is still lying to them. This is what is sometimes referred to as a "bald-faced lie" (cf. Sorensen 2007).⁴⁸

Thomas Carson (2006, p. 290) has an even better example of this sort. A student has been accused of plagiarism. And the student knows that the dean knows that he did it. But the student also knows (based on the dean's reputation) that he will not be punished unless he confesses. So, when the student is called into the dean's office, he

⁴⁶ Philosophers sometimes make a distinction between (a) what a person intends to do by performing an action and (b) what a person foresees as a likely "side effect" of performing that action (cf. Carson 2006, p. 291, Knobe 2006, pp. 208-09). It may be that this distinction does not matter for whether someone is lying.

⁴⁷ Cf. Carson 2006, p. 289. If such a videotape exists, the witness knows that he is likely to be charged with perjury (i.e., lying under oath). But he may very well be more afraid of Tony than he is of a perjury charge.

⁴⁸ Tim Kenyon refers to such statements as "cynical assertions." See his "Cynical Assertion: Convention, Pragmatics, and Saying "Uncle"." *American Philosophical Quarterly*, XL (2003):241-48. He does not count such statements as lies, but only because he assumes that the standard philosophical definition of lying is correct. However, there is empirical evidence that most people are disposed to count such statements as lies. Psychologists at the University of Oregon (Taylor et al. 2003) wanted to see whether young children could distinguish between lying and pretending. For this study, the psychologists wrote several brief stories where the protagonist lies. In order to create pretending stories, they simply modified the lying stories to remove the intent to deceive. For example, they changed the story so that everyone in the story was aware that the protagonist was not telling the truth. In the experiment, the children correctly identified all of the lying stories as involving lying. However, the children "incorrectly" identified many of the pretending stories as involving lying. When they repeated the experiment using adults instead of children, the psychologists got exactly the same results. Since the "pretending stories" did not have the prototypical features of pretending (e.g., taking on a role to have fun), the psychologists had arguably created stories that involve bald-faced lying rather than pretending. When the psychologists subsequently

denies having plagiarized. Although the student does not expect the dean to be deceived, he is pretty clearly lying to the dean.^{49,50}

These cases show that lying is not always about deception.⁵¹ We typically lie in order to deceive other people. And we want these other people to be deceived because that serves our purposes in some way. For example, my new acquaintance will be more impressed with me (or so I believe) if she is deceived about my occupation. But, as the witness and plagiarist cases show, lies can sometimes serve useful purposes even when they do not deceive (cf. Carson 2006, p. 295, Sorensen 2007, p. 262).

So, intending to deceive is not a necessary condition on lying. Thus, we need some other way to rule out statements that are followed with a wink or that are part of a play.

7. Asserting

A number of philosophers (e.g., Chisholm and Feehan 1977, p. 152, Adler 1997, p. 435, Watson 2004, p. 72, Sorensen 2007, p. 256) agree with me that lying is simply

went on to create pretending stories that did have such features, the children and the adults were perfect at distinguishing lying from pretending.

⁴⁹ Kant (1930, p. 227) actually describes a similar case where “my enemy takes me by the throat and asks where I keep my money.” In that case, “the thief knows full well that I will not, if I can help it, tell him the truth.” However, Kant concludes that “my untruth is not a lie because the thief ... has no right to demand [the truth] of me.” In addition, Kant (1959 [1785], p. 40) suggests that, if an assertion is known to be false, we “would only laugh at any such assertion as vain pretense.” In other words, there would be no point in lying in such a case. But there is a point to the student lying about having plagiarized even if he does not expect to deceive anyone.

⁵⁰ This case is a counter-example to **IDL*** as well as **FDL**. The student does not expect to change the dean’s mind about his guilt. And it is not clear that you can *intend* to do something that you do not expect to succeed at (cf. Alfred R. Mele, “Intention, Belief, and Intentional Action,” *American Philosophical Quarterly*, XXVI (1989):19-30, here pp. 19-20). But even if the student could intend to change the dean’s mind, he still might not intend to do so. The student knows that he can get what he wants (viz., getting away with plagiarism) simply by asserting his innocence. So, he may have no interest at all in actually convincing the dean of his innocence.

⁵¹ Even if the speaker does not have to intend to deceive, it might be suggested that someone involved must intend that someone be deceived (cf. Siegler 1966, pp. 129-30). For example, Tony presumably intends

asserting what you believe to be false. Asserting has a normative component that goes beyond merely making a statement.⁵² And it is this normative component that will rule out statements that are followed with a wink or that are part of a play.

While a number of philosophers agree with me that **AL** is the correct definition of lying, only a few (e.g., Chisholm and Feehan 1977, Carson 2006) have tried to specify precisely what the normative component of assertion is.⁵³ In this section, I will argue that these accounts of assertion (unlike my own) yield the wrong results when they are plugged into **AL**.

7.1. Chisholm and Feehan on Asserting

Roderick Chisholm and Thomas Feehan (1977, p. 152) think that whether you have asserted something depends on exactly what you believe. In particular, you *assert* something when (a) you say something and (b) you believe that you are in a situation where your audience is justified in believing (i) that you believe what you say and (ii) that you intend that they believe that you believe what you say. Plugging this account of assertion into **AL**, we get Chisholm and Feehan's definition of lying.⁵⁴

You *lie* to X if and only if: **(CFL)**

1. You state that *p* to X.

that the jury be deceived about his whereabouts. But it is not clear that anyone involved in the plagiarist case intends that anyone be deceived.

⁵² Merely stating something has a normative component. For example, we have to obey many semantic and syntactic rules in order to communicate with each other (cf. Williams 2002, p. 86). But asserting something has a normative component that involves truth (and/or justification) as well as meaning (cf. Chisholm and Feehan 1977, p. 151).

⁵³ Roy Sorensen (2007, pp. 255-56) offers another account of assertion. He claims that an assertion must have "narrow plausibility." In other words, "someone who only had access to the assertion might believe it." (Assertions need not have "wide plausibility." Bald-faced lies, for example, are assertions.) But it is not clear to me why statements that are followed with a wink or that are part of a play do not have such narrow plausibility. In other words, it is not clear to me why the speaker in such cases is only *pretending to assert* rather than *actually asserting* on Sorensen's account.

⁵⁴ Barry O'Neill (2003) has recently proposed a formal definition of lying that is very similar in spirit to Chisholm and Feehan's definition.

2. You believe that, by making this statement, X becomes justified in believing that you believe *p*.
3. You believe that, by making this statement, X becomes justified in believing that you intend X to believe that you believe *p*.
4. You believe that *p* is false.

This definition of lying correctly rules out statements that are followed with a wink or that are part of a play (cf. Chisholm and Feehan 1977, p. 152). In these cases, the speaker does not believe that his audience is justified in believing that he believes what he says.

However, this definition of lying is clearly too narrow. Essentially, under this definition, you are only lying if you expect that you will be successful in deceiving someone about what you believe.⁵⁵ But when I seriously say to my new acquaintance, “I am the Prince of Denmark,” I do not necessarily think that she is justified in believing that I am the Prince of Denmark or that I believe that I am. She may very well be suspicious of strange men at fancy parties who claim to be royalty. Even so, I am clearly lying to her when I make this statement.⁵⁶

However, the Chisholm-Feehan definition can easily be modified to avoid this problem. Instead of requiring that the listeners be justified in believing that the speaker

⁵⁵ More precisely, you expect that, if you end up being unsuccessful in deceiving someone about what you believe, it will be the fault of your listeners, who fail to believe what they are justified in believing.

⁵⁶ In fact, I am still lying to her even if I do not expect her to believe me. Brian Huss (2006, pp. 128-29) concludes that bluffing is not the same as lying on the grounds that you can bluff someone even when she suspects that you might be trying to deceive her. However, you can also lie to someone even when she suspects that you might be trying to deceive her. In such a case, you do not expect your listener to simply believe what you say. Even so, you can still intend to deceive her with your lie. In particular, you may hope to at least increase her degree of belief in a false proposition (cf. Daya Krishna, "Lying' and the Compleat Robot," *British Journal for the Philosophy of Science*, XII (1961):146-49, here p. 147, Chisholm and Feehan 1977, p. 145).

believes what she says, we can simply require that the listeners be *given a reason to believe* that the speaker believes what she says.

You *lie* to X if and only if: (CFL*)

1. You state that *p* to X.
2. You believe that, by making this statement, X has been given a reason to believe that you believe *p*.
3. You believe that, by making this statement, X has been given a reason to believe that you intend X to believe that you believe *p*.
4. You believe that *p* is false.

The fact that someone says something is usually some evidence that he believes it. For example, when I seriously say to my new acquaintance, “I am the Prince of Denmark,” I *do* think that she has been given a reason to believe that I am the Prince of Denmark and that I believe that I am. This evidence that a speaker believes a particular proposition may sometimes be outweighed by other evidence that he does not believe it. But it is still evidence that he believes it.⁵⁷

In some circumstances, however, the fact that someone says something is no evidence at all that they believe it. For example, the fact that an actor on stage says “I am the Prince of Denmark” gives us no reason to believe that he really believes that he is the Prince of Denmark.

In addition, there are *lies* that give us no reason to believe that the liar believes what she says. For example, the student knows that the dean has no reason to believe that the student believes what he says. Also, the student knows that the dean has no reason to

believe that the student intends the dean to believe that the student believes what he says. Everybody involved in this situation knows that everybody is just going through the motions. Nevertheless, the student has clearly *asserted* his innocence and, thus, lied to the dean. Similarly, when someone who is known to be an inveterate liar makes a statement, there is no reason for anyone to believe that she believes that the statement is true. So, if she knows that she is known to be an inveterate liar, the conditions of **CFL**^{*} will not be satisfied. But presumably, someone who is known to be an inveterate liar can still lie.⁵⁸ Thus, **CFL**^{*} is still too narrow.⁵⁹ As Carson (2006, p. 292) puts it, “Chisholm and Feehan’s definition has the very odd and unacceptable result that a person who is notoriously dishonest couldn’t tell lies to those he knows distrust him.”

7.2. Carson on Warranting the Truth

Carson does not think that lying is asserting something that you believe to be false. In fact, Carson (2006, pp. 299-301) explicitly argues against several philosophers who hold this view of lying. Instead, according to Carson (2006, p. 298), (setting aside a few complications) you lie if you *warrant the truth* of something that you believe to be false.⁶⁰ And one warrants the truth of a statement when one makes the statement in a

⁵⁷ According to Richard Foley, we cannot simply equate *evidence* with *reasons for belief*. See his “Evidence and Reasons for Belief,” *Analysis*, LI (1991):98-102. But the few exceptions to this equivalence can be safely ignored for purposes of this paper.

⁵⁸ By contrast, a *pathological liar* (as opposed to someone who merely has a long reputation for lying) may not be a liar at all (cf. Ekman 1997, p. 334). As with the Imperial Stormtrooper, it is not clear that his stating that *p* is an intentional action.

⁵⁹ It would not help to weaken the condition so that it is only required that the listeners be *given what they take to be a reason to believe* that the speaker believes what she says. In the plagiarist case and the inveterate liar case, the statement is not actually evidence that the liar believes what she says *and* it is not taken by the listeners to be evidence.

⁶⁰ One of these small complications will be discussed below in the text. The other two are that, according to Carson, (a) your statement actually has to be false and (b) you only have to fail to believe that it is true. As noted above in section 5, it does not seem that a lie actually has to be false. Also, if a speaker just fails to believe that his statement is true, it sounds to me more like *bullshitting* than lying. According to Frankfurt (2005), bullshitting requires a *lack of concern* for the truth rather than just ignorance of the truth. However, *asserting* something when you do not know whether it is true does indicate a lack of concern for

context where “one promises or guarantees, either explicitly or implicitly, that what one says is true” (Carson 2006, p. 294). Furthermore, it is important to note that whether one has warranted the truth of a statement is independent of what one intends or *believes* (cf. Carson 2006, pp. 296-98).

You *lie* to X if and only if: (CL)

1. You state that *p* to X.
2. You warrant the truth of *p* by making this statement.
3. You believe that *p* is false.

Despite his rejection of **AL**, however, Carson’s definition of lying is equivalent to saying that (a) lying is asserting something that you believe to be false and (b) warranting the truth is the normative component of assertion (cf. Carson 2006, p. 300). In fact, several philosophers have characterized the normative component of assertion in exactly this way.⁶¹ According to these philosophers, whenever we make an assertion, we promise or guarantee to our audience that what we say is true. In addition, our audience has the right to rebuke us if our claim turns out to be false. Thus, when we assert *p*, we take on a certain responsibility with respect to *p*. We do not take on the responsibility of actually making *p* true (cf. Carson 2006, p. 294). But we do make a “commitment to the defensibility of *p*” (Watson 2004, p. 68).

CL gets right the cases that we have considered so far. For example, it rules out statements that are followed with a wink or that are part of a play. In these cases, the speaker does not offer a guarantee that what she says is true. But it rules in the

the truth. (Of course, as Frankfurt (2002, p. 341) points out, bullshitting and lying are not necessarily mutually exclusive categories.) But in any event, my definition can easily be modified to incorporate both of these complications. As I discuss in the text, my main objections to Carson’s definition lie elsewhere.

statements made by the plagiarist and the inveterate liar. In these cases, the speaker clearly warrants the truth of what she says.

As Carson (2006, pp. 297-98) himself recognizes, however, **CL** is too broad as it stands. For example, suppose that a politician has been asked to give a serious speech at one banquet and a humorous speech at another banquet. Further suppose, however, that the politician gets his dates mixed up and accidentally delivers the humorous speech to the audience expecting a serious speech. The humorous speech includes a joke about the President having “broken wind” at a meeting with foreign dignitaries (an event which did not actually occur). Carson’s intuition is that the politician is not lying in this case even though he says something that he believes to be false. However, since this is a context where his audience expects a serious speech, the politician (according Carson’s definition) warrants the truth of his statements.

In order to rule out such cases, Carson (2006, p. 298) has added an additional condition to his definition. According his revised definition, the politician is not lying because he *believes* that he is not warranting the truth of his statements.

You *lie* to X if and only if: (CL^{*})

1. You state that *p* to X.
2. You warrant the truth of *p* by making this statement.
3. You do not believe that you are not warranting the truth of *p* by making this statement.
4. You believe that *p* is false.

⁶¹ See, e.g., Charles Sanders Peirce, "Belief and Judgment," in Charles Hartshorne and Paul Weiss, eds., *Collected Papers*, vol. V, (Cambridge: Harvard, 1934), pp. 376-87, Brandom 1983, Watson 2004.

However, **CL**^{*} is still too *narrow*. This can be shown by considering a simple variation on Carson's own example. Suppose that, when the politician gets his dates mixed up, he accidentally delivers the serious speech to the audience expecting a humorous speech. Given that he believes that his audience expects a serious speech, my strong intuition is that the politician *is* lying in this case if he makes statements that he believes to be false.⁶² However, since this is a context where his audience actually expects a humorous speech, the politician does not (according Carson's definition) warrant the truth of his statements.⁶³ Thus, according to Carson's definition of lying, the politician is not lying.

7.3. Believing that One is Warranting the Truth

According to Carson's definition of warranting, whether a speaker has warranted the truth of a statement is independent of what the speaker intends or *believes*. However, it seems that his beliefs do matter to whether he is lying.⁶⁴ In both versions of the mixed-up politician case, for example, what the politician *believed* that he was doing seemed to be a critical factor in our judgment about whether he was lying. This is right in line with Saint Augustine's claim that "a person is to be judged as lying or not lying according to the intention of his own mind, not according to the truth or falsity of the matter itself."⁶⁵

⁶² We actually do not have to rely on intuitions about cases where a speaker unknowingly fails to warrant the truth of a statement in order to see that warranting the truth of a statement is not necessary for lying. In section 7.3 below, I give a more mundane counter-example to **CL**^{*} where a speaker lies despite explicitly telling his audience that he does not warrant the truth of his statement.

⁶³ The politician *believes* that he is warranting the truth of what he says. But he is wrong (cf. Carson 2006, p. 297).

⁶⁴ This is certainly suggested by the condition that Carson added to his definition. While some speech acts (such as *saying* something) can be performed without knowing that you are doing it, other speech acts (such as *promising* something) require that you know that you are doing it. See David Owens, "A Simple Theory of Promising," *Philosophical Review*, CXV (2006):51-77, here p. 54. Lying seems to fall into this second category as well.

⁶⁵ Augustine is talking specifically about what the speaker believes about the statement itself. But it seems that the point applies more generally. Although it is ultimately unsuccessful, a virtue of Chisholm and

Thus, instead of saying that a speaker has to actually warrant the truth in order to assert something, we might just say that a speaker has to *believe* that he is warranting the truth.

You *lie* to X if and only if: (BWL)

1. You state that *p* to X.
2. You believe that you warrant the truth of *p* by making this statement.
3. You believe that *p* is false.

BWL gets both versions of the mixed-up politician case correct. But unfortunately, this definition still rules out some cases of lying that should not be ruled out. For example, suppose that the witness follows up his statement that “Tony was with me at the time of the murder” by saying, “Of course, you know I am really bad with dates and times.” This proviso makes it clear that the witness is not *guaranteeing* that Tony was with him at the time of the murder.⁶⁶ Thus, he is not (according to Carson’s definition) warranting the truth of what he says.⁶⁷ In addition, he clearly *does not believe* that he is warranting the truth of what he says. Thus, he is not (according to **BWL**) lying. However, if he believes that Tony was not with him at the time of the murder, it still seems pretty clear that he has lied to the jury. The witness is not taking any responsibility for the truth (or defensibility) of his statement about Tony’s whereabouts at the time of the murder. (If the jury decides to rely on his statement and it turns out to be false, the responsibility lies with them.) But it still seems that he intends to be taken seriously as having expressed what he believes to be the case. In other words, he is still making an

Feehan’s definition is that it does cash things out solely in terms of what the speaker believes. Peter Graham (1997) makes an analogous point with respect to defining *testimony*.

⁶⁶ The witness may be lying about being bad with dates and times (as well as about Tony’s whereabouts). But regardless of whether he actually is bad with dates and times, his merely *saying* that he is bad with dates and times means that he is not warranting the truth of his statement about Tony’s whereabouts. (Thanks to Kay Mathiesen for this example.)

assertion about Tony's whereabouts at the time of the murder (cf. Sorensen 2007, p. 255). Thus, this case shows that you can be lying (when you say something that you believe to be false) even if you explicitly say that you are not warranting the truth of what you say.⁶⁸

8. Some Virtues of My Definition

My definition correctly handles the counter-examples to the other proposed definitions of lying. The plagiarist, the inveterate liar, the politician giving a serious speech, and the witness who claims to be bad with dates and times all believe that they are in a situation where they should not say what they believe to be false. Thus, according to **BNL**, they are lying when they say something that they believe to be false.

In addition, my definition is not open to the sort of counter-example considered in the previous section. For example, suppose that the witness follows up his statement that "Tony was with me at the time of the murder" by saying, "Of course, I think that it is ok to say things that I believe are false in this sort of situation." Following his statement with this proviso is just like following his statement with a wink (to indicate that he is not to be taken seriously). Thus, this proviso makes it clear that Grice's first maxim of quality is not currently in effect. In addition, he clearly *does not believe* that Grice's first maxim of quality is currently in effect. Thus, he is not (according to **BNL**) lying. But this seems like exactly the right result in this case. He cannot explicitly say that he does not take himself to be subject to Grice's first maxim of quality and still be making an

⁶⁷ This case is a counter-example to **CL*** as well as **BWL**. The witness is lying even though he is not actually warranting the truth of his statement.

⁶⁸ There is a counter-example to Gary Watson's (2004, pp. 67-70) account of the norm of assertion along the same lines. In that case, the relevant proviso is "Of course, I am not committed to the defensibility of what I just said." Admittedly, this counter-example is not quite as felicitous sounding as my counter-example to **BWL**. But it still seems to work. And my conjecture is that any account of the norm of assertion other than my own can be handled in the same sort of way.

assertion. In fact, he is essentially saying with this proviso that he is not making an assertion. Thus, he does not seem to be lying when he says something that he believes to be false about Tony's whereabouts at the time of the murder.

Another virtue of my definition is that it provides a straightforward analysis of Augustine's (1952, p. 57) example of an *altruistic lie*.⁶⁹ A man tells his friend that there are no bandits on a certain road even though he believes that there are. He does this because he knows that his friend does not trust him and will conclude from his statement that there *are* bandits on the road.⁷⁰ This person is lying under my definition (**BNL**), because he says something that he believes to be false and he believes that Grice's first maxim of quality is in effect. But, surprisingly enough, he is not lying under the standard philosophical definition of lying (**IDL***). This person does not intend to deceive his friend about there being no bandits on the road.⁷¹ He wants his friend to correctly believe that there *are* bandits on the road. In fact, this person does not even intend to deceive his friend with respect to his believing that there are no bandits on the road.⁷²

9. Some Objections to My Definition

I am a big fan of the television show *Homicide*. Regularly on this show, the homicide detectives put suspects in "the box" and try to get them to confess. A standard

⁶⁹ The scientist that fakes her data in order to trick people into believing a theory that she believes to be true is another example of an altruistic lie (cf. O'Brien 2007, p. 228). This example counts as a lie on both my definition and the standard philosophical definition.

⁷⁰ In the same passage, Augustine also gives an example of a *deceptive truth*. A man truthfully tells his enemy that there are bandits on a certain road. He does this because he knows that his enemy does not trust him and will conclude from his statement that there are no bandits on the road. Although he does intend to deceive his enemy with his statement, this person is not lying on both my definition and the standard philosophical definition.

⁷¹ This person does intend to deceive his friend. In particular, he intends that his friend believe that he intends that his friend believe that there are no bandits on the road. Thus, this case would count as a lie under **IDL**. But, as noted above in section 6, **IDL** is too broad as a definition of lying.

technique for achieving this goal is to lie to the suspect.⁷³ For example, the suspect is told that his fingerprints have been found on the murder weapon, that his DNA has been found at the crime scene, or that his partner has just given him up.

This sort of case might seem to be a counter-example to my definition of lying. The cops are pretty clearly lying to the suspect. But there does not seem to be a norm in effect that the cops should not say what they believe to be false (quite the contrary, in fact). If that were true, they would not be lying according to my definition. Thus, my definition would be too narrow. However, I contend that the norm *is* in effect in “the box” and that the cops believe that it is.⁷⁴

The cops believe that they are in a situation where they should not say things that they believe to be false. It is just that other interests of the cops (namely, getting the suspect to confess) override this norm (cf. Simon 1991, pp. 211-12).⁷⁵ In other words, the norm is in effect and the cops have violated it, but they have a good *excuse* for doing so.⁷⁶ Similarly, the posted speed limit is still 35 miles per hour even if I decide that it is

⁷² Some definitions require an intention to deceive with respect to your believing *p* in addition to (or instead of) an intention to deceive with respect to *p* itself (cf. Mahon 2008, section 1.6).

⁷³ Cf. David Simon, *Homicide* (New York: Ivy Books, 1991), pp. 216-17, Christopher Slobogin, "Deceit, Pretext, and Trickery: Investigative Lies by the Police," *Oregon Law Review*, LXXVI (1997):775-816, here pp. 785-88.

⁷⁴ The cops are clearly making use of the fact that at least the suspect believes that the norm is in effect.

⁷⁵ The norm is *explicitly* in effect for a witness who is testifying under oath, but the witness may have interests (e.g., avoiding the wrath of Tony and his crew) that outweigh the force of the norm. A *white lie* is another example where Grice's first maxim of quality is in effect, but is overridden by other interests (e.g., not offending someone) that we have. In addition to coming into conflict with other interests that we have, the norm in question ("Do not say what you believe to be false") can come into conflict with other norms. For example, it can come into conflict with the moral norm "One ought to prevent harm to others" (à la W. D. Ross, *The Right and the Good* (Oxford: Oxford, 1930)) as when the only way to save a friend from a murderer is to lie (cf. Kant 1949 [1797]). But my claim is that, regardless of what else is going on, disobeying this norm when you believe that it is in effect is lying.

⁷⁶ Norms of conversation can "clash" with each other as well as with other interests that we have (Grice 1989, p. 30). For example when a teacher who believes in creationism has to give a lesson on evolution, Grice's first maxim of quality comes into conflict with Grice's second maxim of quality. If the teacher only violates the norm against saying what she believes to be false in order to obey the norm against saying that for which she lacks adequate evidence, some (e.g., Lackey 2007, p. 602) might want to say that she is

best to rush my injured friend to the hospital at 70 miles per hour. I have broken the speed limit even though I may have a good excuse for doing so.

Everyone involved in the interrogation seems to recognize that the norm is in effect. For example, after he learns of the deception, the suspect will often complain that he has been lied to. In addition, while the cops rarely have much sympathy for the suspect, they do seem to recognize that it is reasonable for him to complain. By contrast, the norm is not in effect at all if you are performing a play or if you wink to indicate that you are not to be taken seriously. If a member of the audience complained that she had been lied to, she would have to be crazy and the cast would undoubtedly think that she was crazy.⁷⁷

I am also a big fan of *Star Wars* and my favorite character is the pilot of the *Millennium Falcon*, Han Solo. In one memorable scene, our heroes use the garbage chute to escape from the detention block of the *Death Star*. But once they land in the garbage, Han Solo sarcastically says, “The garbage chute was a really wonderful idea. What an incredible smell you’ve discovered!”

This sort of case might also seem to be a counter-example to my definition of lying. Solo is pretty clearly not lying.⁷⁸ Just the opposite, in fact. He is trying to communicate something that he believes to be true (viz., that the garbage chute was a

not lying. In order to accommodate that intuition, my definition might be modified to include an exemption for such cases.

⁷⁷ There are other contexts in which it is ok to violate the norm, but not simply because the speaker has a good excuse. For example, the *Liar’s Club* was a television game show on which a panel of celebrity guests passed around some unusual object and gave bizarre explanations of its purpose. The contestants then had to guess which of the celebrity guests was *not* lying to them. The celebrity guests had essentially been given permission to violate the norm in this context. (It is perfectly ok to bluff in poker for the same reason.) Thus, even though they have also been lied to, it would be crazy for contestants on this show to complain about it in the way that it is not crazy for the suspect to complain about it.

⁷⁸ Han Solo is not above telling lies (e.g., about the speed of the *Millennium Falcon*). But he is not lying here. I would like to thank Daniel Silvermint for this example of conversational implicature. I would also

really *bad* idea). But he is certainly *saying* something that he believes to be false. In addition, 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. However, 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.

Solo's comment about the garbage chute is an example of *conversational implicature*.⁷⁹ Solo "blatantly fails to fulfill" or "flouts" the first maxim of quality in order to express something other than what he literally says (Grice 1989, p. 30). I contend that, by flouting this norm of conversation, Solo turns it off.

The default whenever we engage in linguistic behavior is that the first maxim of quality (as well as the rest of Grice's maxims) is in effect. However, these norms can be turned off or suspended. Basically, you can signal to your audience—in a manner that is collectively accepted by the relevant linguistic community for this purpose—that you will not be obeying the norm of conversation in question.⁸⁰ If you give such a signal, other people will not expect you to obey the norm, they will not think that you ought to obey the norm, they will not think that you are subject to rebuke for not obeying the norm, etc. Winking is an example of a signal that turns off Grice's first maxim of quality in many cultures.

like to thank Marc Moffett for pressing this objection in his excellent comments on this paper at the 2008 Pacific Division Meeting of the American Philosophical Association.

⁷⁹ Cf. Grice 1989, pp. 24-37, Wayne Davis, "Implicature," *Stanford Encyclopedia of Philosophy*, (2005): <http://plato.stanford.edu/entries/implicature/>.

⁸⁰ In other words, there is a *norm* that a particular signal turns off the norm of conversation in question. Of course, even if they do not believe that she intended to send such a signal, it will sometimes be clear to an audience that a speaker is not to be taken seriously. For example, if a witness in court says with a straight face that she is a Giant Panda, the best explanation may be that she is crazy. Now, if she really is crazy, but believes that she is actually a Giant Manatee, then she is clearly lying. But if she is not crazy and is only trying to convince everyone that she is crazy or is just trying to disrupt the proceedings, some (e.g., Mahon 2008, section 1.6) might want to say that she is not lying. In order to accommodate that intuition, it might

While a collectively accepted signal is often an overt behavior such as a wink, it can also be quite subtle. In fact, it may just be clear from the context that the speaker does not really mean to seriously assert what he is actually saying. For example, the deadpan humorist Steven Wright is not lying when he says with a completely straight face, “I can levitate birds. No one cares.”⁸¹ Similarly, when a speaker flouts Grice’s first maxim of quality as Han Solo did, the best explanation of the utterance is that the speaker means to express something other than what he is actually saying.⁸²

A norm of conversation does have to be in effect at the time of utterance in order for conversational implicature to work. For example, the fact that Grice’s first maxim of quality is in effect (and the fact that Solo is clearly not acting in accordance with it) is part of what allows Princess Leia to figure out what Solo really means to express. But in the case of both winking and flouting, the norm *is* in effect when the utterance begins. The signal essentially has the effect of retroactively turning off the norm.⁸³ For example, we usually wink *after* we make a statement that we do not want to be taken seriously.⁸⁴

be suggested that pretending to be crazy makes it reasonable for one to believe that the norm is not in effect (even though it does not make it common knowledge that the norm is not in effect as winking would).

⁸¹ It is not just the absurdity of the content that provides the indication. A completely preposterous statement can still be a (bald-faced) lie (cf. Sorensen 2007, p. 253).

⁸² Davis (2005, section 8), for example, suggests that people use *inference to the best explanation* to work out conversational implicatures.

⁸³ It should also be noted that the signal typically turns off the norm with respect to that one single utterance. However, some might object that, instead of actually turning off the norm, flouting just makes disobeying the norm permissible. In that case, another condition would simply need to be added to my definition of lying. Namely, you must not believe that you have signaled to X—in a manner that is collectively accepted by the relevant linguistic community for this purpose—that you mean to express something other than *p*.

⁸⁴ Interestingly, it seems that the signal can be sent fairly long after the statement is made and still turn off Grice’s maxim with respect to that statement. For example, some people will not give any indication that they are joking until they have you completely convinced that they are serious. But such “leg pulling” does not seem to count as lying.

By contrast, a bald-faced liar, such as the plagiarist, gives no such signal. Although he blatantly fails to fulfill the norm, it is still in effect.⁸⁵ The best explanation of his utterance is that he means to express exactly what he says. So, unlike Solo, he is lying.

10. Deceptive Lying

While lying does not require an intention to deceive, it must be acknowledged that philosophers are primarily interested in lies that are intended to deceive.⁸⁶ For example, Kant is clearly concerned with cases of lying that have the potential to destroy trust. He claimed that it is always wrong to lie because lying whenever it is to your advantage is a self-defeating moral rule. If everyone followed this rule, we could not trust anybody when they spoke and there would be no point in lying.⁸⁷ Thus, it is not a problem for Kant if things like the plagiarist case are ruled out by a definition of lying (cf. Sorensen 2007, p. 263). There is certainly something morally problematic about what the student does in the plagiarist case.⁸⁸ But it is not clear that it destroys (or even decreases) trust. As Roy Sorensen (2007, p. 252) puts it, “bald-faced lies do not fool anyone. They are no more a threat to truth telling than sarcastic remarks.”

In addition, when epistemologists talk about lying, they are principally concerned with *deceptive* testimony that can interfere with people’s ability to acquire knowledge.

⁸⁵ Grice (1989, p. 30) seems to suggest that you can only “quietly and unostentatiously *violate* a maxim.” But there is certainly a sense in which the plagiarist “blatantly” violates the first maxim of quality.

⁸⁶ The very title of Chisholm and Feehan’s (1977) paper suggests that this is the type of lying that they were specifically interested in.

⁸⁷ I do not take any position on whether Kant’s argument is correct. (In fact, Elliott Sober argues that, while a world in which everybody lies is unstable, it is not impossible. See his “The Primacy of Truth-Telling and the Evolution of Lying,” in *From a Biological Point of View*, (Cambridge: Cambridge, 1994), pp. 71-92.) I just point out that the argument requires a type of lying that involves the intent to deceive.

⁸⁸ Thus, as Sorensen (2007) points out, bald-faced lies clearly have some philosophical interest.

Thus, it may be no problem for epistemologists if things like the plagiarist case are ruled out by a definition of lying. There is nothing epistemically problematic at all about what the student does in the plagiarist case. In particular, it is not going to lead the dean to acquire false beliefs.

Since we are looking for a definition of lying for *philosophical* use (albeit one that comes close to common usage), perhaps our definition should include an “intent to deceive” condition after all. Thus, it might be suggested that the standard philosophical definition (**IDL**^{*}) is a perfectly good stipulative definition for our purposes. But although **IDL**^{*} would be a simpler definition of *deceptive lying*, I would like to suggest that it would actually be better to add an “intent to deceive” condition to my definition of lying rather than revert to the standard philosophical definition.

We would presumably like a unified account of lying. In particular, deceptive lying should be a proper subset of lying in general. However, **IDL**^{*} seems to rule in cases of verbal deception that are not lies at all. For example, an actor on stage is not lying when he recites a line that he happens to believe to be false. But it is possible that the performance is part of an elaborate deception aimed at getting members of the audience to believe that that particular line from the play is actually true. In other words, the actor says something that he believes to be false with the intent to deceive people about that thing.⁸⁹ By contrast, my proposal correctly rules out such cases because Grice’s first maxim of quality is not in effect in this context.⁹⁰

⁸⁹ In order to satisfy the first condition of **IDL**^{*}, it would also have to be a line that the actor addresses *to the audience* of the play (e.g., in an aside) rather than to another actor. As with this case of the deceptive actor, it is not clear whether an adult telling a tall tale to credulous children is actually lying to them even if he intends to deceive them and, thus, satisfies the conditions of **IDL**^{*} (cf. footnote 21 above).

⁹⁰ There may be one possible role for the standard philosophical definition. It might turn out to be the case that young children learn how to deceive with words before they learn (implicitly) when the norms of conversation are in effect. If so, we might want to refer to some of their (pre-norm) utterances as lies.

11. Conclusion

According to the standard philosophical definition, lying is saying something that you believe to be false *with the intent to deceive*. However, people sometimes “go on the record” with something that they believe to be false even though they have no intent to deceive. In response to such examples of non-deceptive lying, it has been suggested that lying is simply *asserting* something that you believe to be false (where asserting something involves a normative component that goes beyond merely saying something). However, the existing accounts of asserting something give us a definition of lying that still does not capture common usage. In this paper, I have argued that the correct definition of lying is that (a) you say something that you believe to be false and (b) you believe that you are in a situation where the following norm of conversation is in effect: “*Do not say what you believe to be false.*” But I do note that philosophers typically *are* interested in lies that are intended to deceive. So, for most philosophical purposes, I propose a definition of *deceptive lying* which simply adds an “intent to deceive” condition to my definition of lying.⁹¹

Don Fallis
School of Information Resources
University of Arizona

⁹¹ I would like to thank Jonathan Adler, Jody Azzouni, Matt Bedke, Tom Carson, Jane Carter, Jacob Caton, David Christensen, Tony Doyle, Martin Frické, Jerry Gaus, Allen Habib, Terry Horgan, Peter Lewis, Pen Maddy, James Mahon, Kay Mathiesen, Marc Moffett, Shaun Nichols, Patrick Rysiew, Roy Sorensen, Amie Thomasson, Susan Vineberg, Dennis Whitcomb, and Dan Zelinski for much helpful feedback. I would also like to thank audiences at the Philosophy Department at the University of Arizona, at the Philosophy Department at Wayne State University, at the Department of Logic and Philosophy of Science at the University of California, Irvine, at the Philosophy Club at Rutgers University, at the Philosophy Club at the John Jay College of Criminal Justice, and at the Pacific Division Meeting of the American Philosophical Association.