

# Lying and Denying

---

Ryan Doody

rdood@mit.edu

## 1 Introduction

It's natural to distinguish between two importantly different ways one can be deceptive in linguistic communication: one can deceive you by *lying to you*, or one can deceive you by *deliberately misleading you*. The distinction is standard, but elusive. What exactly *is* the difference between being *lied to* and being *deliberately misled*? Just as standard as the distinction itself, is the notion that (*ceteris paribus*) there is something worse, morally, about lying than misleading.

**Here's an example.**<sup>1</sup> You are dangerously allergic to peanuts --- the slightest waft of peanut-matter will send you into anaphylactic shock. And we both know this. I, unbeknownst to you however, very much want you dead. I have invited you over for dinner. I cooked all of your food in *peanut oil*. You've learned to be extra careful about your food allergies. The following exchange occurs.

**Case 1.** (1) **You:** Are there any peanuts in the food?

**Me:** No. The food is safe for you to eat.

In this version of the story, it seems fairly clear that I have lied to you. I know that your food is *not* safe for you to eat. But I explicitly told you it was.

**Case 2.** (2) **You:** Are there any peanuts in the food?

**Me:** No. There aren't any peanuts in the food.

In this version of the story, it is *less* clear that I have lied to you. The food was prepared with *peanut products*; there are no whole peanuts in the food. Nevertheless, my answer was deliberately crafted to suggest to you that the food is safe for you to eat. I have been *misleading*. But I have not *lied*.

---

<sup>1</sup>This is lifted from Saul's book *Lying, Misleading, and What Is Said*.

This distinction should feel familiar. In the first version of the story, I have lied to you. In the second version of the story, I have done something *deceptive* --- but I haven't lied. What's the difference?

I argue that the distinction is slipperier, and a lot more *pragmatic*, than is often assumed. Ultimately, the distinction between the two acts amounts to a distinction between what we can, in the case of misleading, and cannot, in the case of lying, plausibly get away with. This account of the distinction heavily undermines the widespread ethical belief that, as deceptions go, lying is *morally worse* than misleading. (I think, however, that there are independent reasons to be suspicious of this moral distinction). In response, I present a speculative story about how, and why, it is that creatures like us could come to deeply internalize a moral preference for mere misleading over outright lying --- even if we ultimately lack good moral reasons for having such a preference.

## 2 *Lying Belongs to Semantics, Misleading to Pragmatics*

What is the difference between lying and misleading? Here's a plausible suggestion. The Lying/Misleading Distinction corresponds, roughly, to the distinction between, on the one hand, *What Is Said* and, on the other, *What Is Communicated* by an utterance. Lying is semantic, merely misleading is pragmatic. Utterances are made in the presence of a conversational context. And features of this background context can make it so that *What Is Communicated* by an utterance is something that goes beyond what the utterance strictly and literally *says*.

Whether or not this suggestion seems right to you is going to depend on how you think the line between Semantics and Pragmatics should be drawn. And how you think *that* line should be drawn will depend on the theoretical role(s) you intend the Semantic/Pragmatic distinction to play. Rather than understand this suggestion (that the Lying/Misleading Distinction corresponds to the Semantic/Pragmatic Distinction) in terms of the various already-established ways of drawing the line between *What Is Said* and *What Is Communicated*, we can, instead, understand it as an organizing idea: the difference between lying and misleading, whatever it is, *does* cor-

respond to the distinction between *What Is Said* and *What Is Communicated* --- at least in some interesting sense. The task then becomes to say what must be the case about the Said/Communicated Distinction, understood in this way, in order for it play the role we want it to. (At the end of day, there might very well be *many* interesting notions of, e.g., *What Is Said* --- each corresponding to a different theoretical role. Explaining the intuitive difference between lying and misleading being one such role).<sup>2</sup>

**Organizing Idea:** The difference between *lying* and *misleading* is that, in the former, but not the latter, the deceptive proposition that's communicated to the audience is part of *What Is Said* by the utterance.

There are various theoretical roles we might want the notion of *What Is Said* to play. For example, we might think it is the content which is attributed to speakers by indirect speech reports (e.g., "Well, *she* said that . . ."). Or, for example, we might think it is the content of whatever is *explicitly* (either consciously or unconsciously) *represented* in the mind of the speaker / audience upon uttering / hearing. Or, we might take it to be whatever content is necessary to count as a minimally competent speaker of the language. Or, we might take it to be the content of *What Is communicated* that is calculated by the domain-specific reasoning of a linguistic faculty. Etc. And here now is another theoretical role: whatever it is that corresponds to the difference between lying and misleading.<sup>3</sup>

---

<sup>2</sup>This is how Saul casts her project in *Lying, Misleading, and What Is Said*. The fact that so-and-so's account of *What Is Said* cannot make sense of the Lying/Misleading Distinction is not, in itself, a good reason to reject so-and-so's account unequivocally --- it only provides a good reason to reject it as an account of the particular version of *What Is Said* that's supposed to be operative in making the distinction between lying and misleading.

<sup>3</sup>I am, mostly in order to simplify the discussion, going to go along with Saul by accepting the Organizing Idea. My acceptance of approach things this way, however, should be understood to be *pro tem*. I am uncomfortable just resting happily with some kind of Semantic Pluralism. It would be good, at the very least, if a lot more could be said about these different theoretical roles, how they could potentially bear on each other, etc. And it would be, in my opinion, very good indeed if they could, somehow, all be brought under the welcoming umbrella of some grand unifying account of communication.

## 2.1 What It Is To Lie, What It Is To Mislead

Together, let's you and I (at least *pro tem*) accept that the Lying/Misleading Distinction corresponds to the distinction between *What Is Said* and *What Is Communicated*. How are the two distinctions related?

Allow me to make the correspondence between the two distinctions significantly more explicit by offering a simplified, very rough (rough to the point of wrong) account of what it is for an utterance in a conversational context to be a *lie*.

An utterance  $u$  by a Speaker  $S$  in a conversational context  $C$  is a *lie* only if

- (a) By uttering  $u$ ,  $S$  says that  $p$
- (b)  $S$  believes  $p$  to be false
- (c)  $S$  intends the uttering of  $u$  to deceive the audience
- (d) The conversational context  $C$  is of the appropriate kind

This account of lying is more-or-less the same as the one in Saul 2012.<sup>4</sup> The nitty-gritty details don't matter. Getting even a dangerously oversim-

---

<sup>4</sup>What are the difference? First, Saul's account omits clause (c). Why? She is concerned about some of the cases described in Sorensen 1997 --- which involve examples of 'bald-faced lies': utterances that say (at least in normal circumstances) something that the falsity of which is taken to be common ground in the conversational context. Given that the speaker knows that everyone knows (etc., etc., ever onward toward infinity) that  $p$  is false, it seems (at the very least) odd to ascribe to the speaker the intention to deceive. These are weird cases. (Frankly, I have trouble making sense of what proposition the assertion of a 'bald-faced lie' purports to communicate at all! I suspect that speech acts like this (the examples of which come from court cases, fascist states, etc.) aren't properly lies. The utterer is doing something non-standard with words). I invite you to join me in ignoring these potential complications. Second, Saul requires that the Speaker take herself to be in a 'warranting context'. This clause prevents us from having to say that, for example, all actors are habitual, professional liars. The stage and screen are not 'warranting contexts'. Clause (d) is intended to get across the same idea. Finally, Saul restricts her account to cases in which  $S$  is "not the victim of linguistic error/malapropism, or using metaphor, hyperbole, or irony" etc. On the one hand, this is unfortunate. So many of the most interesting cases of linguistic deception involve metaphor, etc. On the other hand, the restriction is probably necessary --- if we're to make any progress. There would be far too much up for grabs, otherwise. That being said, I hope to say something about the Lying/Misleading Distinction which will ultimately bear on how we should think about these harder (and, in my opinion, more interesting) cases.

plified account, like this one, on the table helps to illuminate the relationship between *lying* and *saying*.

Look at clause (a). It says that a necessary condition of utterance  $u$  being a lie is that in virtue of uttering  $u$  in that conversational context, the speaker thereby *says* something  $s$ /he believes to be false (or knows to be lacking the appropriate reasons to believe). If the utterance of  $u$  merely *communicates* something the speaker believes to be false (or knows to be lacking the appropriate reasons to believe), that's not enough for  $u$  to count as a lie. In such a case, to utter  $u$  would be to *mislead* (perhaps with the intent to deceive), not to *lie*.

Let's call the proposition  $p$  about which the speaker (b) believes to be false (or knowingly lacks good reason to believe) and (c) intends to get the audience to accept the *deceptive content*. The connection, then, between the two distinctions is this. If the deceptive content is part of *What Is Said*, then the speaker has lied. If the deceptive content is part of *What Is Communicated* but not part of *What Is Said*, the speaker has said something misleading but has *not* lied.

## 2.2 Context and *What Is Said*

The account of the Lying/Misleading Distinction hand-wavingly gestured at above places constraints on how to understand *What Is Said*. Saul 2012 includes an extended discussion of the ways in which the various standard ways of drawing the Semantic/Pragmatic Distinction are not up to task.<sup>5</sup>

---

<sup>5</sup>Saul helpfully sorts the myriad of positions into three categories: Unconstrained Views, Constrained Views, and Austere Views.

**Unconstrained Views** (like, Cappelen & Lepore's *Speech Act Pluralism*) hold, very roughly, that if there exists a possible context in which it would be right to say of a speaker's utterance  $u$  that  $s$ /he *said that*  $p$  then  $p$  can be counted as (perhaps not uniquely) *what is said* by the utterance of  $u$ . The problem? These views will *overclassify* instances of deception as lies.

**Constrained Views** (like, Relevance Theorist's *Explicature*; King, Stanley, Szabo's and Stainton/Taylor's *Semantic Content*; Recanati's *What Is Said*, etc.) hold that there is a tighter connection between the actual sentence uttered and what the sentence says --- but the connection is loose enough to allow conversational context to have some affect on what the utterance says. What features of the conversational context help determine what an utterance says? And in what way? Different views tell different story, but there is general agreement that the contribution context makes in determining *What Is Said* is largely a

### Quick Aside: can a lie be true?

There's a worry that clause (b) is too weak. It should be revised to require not only that *S* believe *p* to be false, but also that *S truly* believe *p* to be false. You cannot truly lie if you're telling the truth (even if only by accident). Or so the worry goes.

I am inclined, however, to think that clause (b) is in fact too *strong*. Less is required to lie than (b) allows. Here's an example.

**Roger's Coin Flip.** Roger is going to flip his lucky fair coin. He will flip it in the privacy of his office, and the report the result to the rest of us. It is going to be quite the event! Sadly for you, you cannot make it. Later, you bump into me on the street. You falsely (but reasonably) believe that I had been there for the big reveal. (Unbeknownst to you, I wasn't there either). You ask me, "How did Roger's coin land?" I respond, confidently, "It landed heads." As it happens, the coin in fact did land heads.

It seems fairly clear to me that I've lied to you --- even though what I've said happens to be true. The problem is that I have *no good reason to think* the coin landed heads. I wasn't there. Moreover, I know the coin is fair. So I also don't believe it's *false* that the coin landed heads either. In order for a lie to be a lie, it needn't be false and, contra clause (b), the liar needn't believe it to be false.

You might think, in fact, that the liar needn't properly *disbelieve* the lie in order for it count as *a lie* even! Amend the above case so that we all know that Roger's lucky coin is, say, biased 90% toward heads. So I am fairly confident, despite not being there for the big reveal, that the coin did in fact land heads. Still, I lied to you. (This perhaps provides some support for the view that knowledge, and not just mere true belief, is the norm of assertion.)

There is potentially a puzzle here. Even if I know antecedently that the coin was likely to land heads, if I wasn't there to witness it, my assertion is a lie. Suppose that I was there for the big reveal. And that, due to the effects of the unrelenting ravages of time on his eyesight, Roger's ability to determine the outcomes of coin toss is compromised --- making him only 90% reliable at accurately reporting such things. If my confidence that the coin landed heads is based on Roger's testimony, then my assertion is *not* a lie. (Or so it seems to me). What's the difference?

Rather than rehearse all of her moves, allow me to draw some very general lessons from her discussion.

First, in order for it do the work we want, *What Is Said* must be truth-evaluable. In particular, it must be the kind of thing about which the audience can come to believe to be true, and the speaker can believe to be false (or know to lack the appropriate reasons to believe to be true).

Second, uncontroversially, *at least sometimes* context is required in order for an utterance to be truth-evaluable.

- (3) I really like Alan's new sweatshirt.
  - a. Ryan really likes Alan's new sweatshirt.
  - b. ? { . . . } really likes Alan's new sweatshirt.
  - c.  $\llbracket I \rrbracket =$  the Speaker

Sentences, like (3), that contain *indexicals* are paradigmatic examples of this. Absent any context at all (3) doesn't seem to express a truth-evaluable proposition at all. Rather, it can be thought of as a function from a context to a truth-evaluable proposition. (A fairly straightforward rule-governed function, but an unsaturated function nevertheless). The naïve hope that *What Is Said* by an utterance can be cloistered-off entirely from the context in which it is uttered is dashed. The semantic content of an utterance cannot always be read off of the context-independent linguistic mean-

---

function of the intentions of the speaker, the expectations of the audience, and the interplay between the two. The Problem? These views place too much emphasis on what the speaker *intends* and what the audience *understands*. And situations in which a deliberate effort is made to be deceptive are situations in which, by their very nature, what the speaker intends the audience to understand comes apart from what the speaker actually intends. If *What Is Said* is tied too closely to what the speaker intends the audience to understand &c., then it will be nearly impossible to deceive without lying.

**Austere Views** (like, Bach's and Borg's, and Cappelen & Lepore's *Semantic Content*) hold that *What Is Said* is \*very\* tightly linked to the uttered sentence. Context makes only very minimal, if any, contribution to what an utterance says. (on some Austere Views, like Bach's, the semantic content of an utterance needn't even be a truth-evaluable proposition. Sentences with indexicals, demonstratives, etc., express proposition-radicals). The Problem? There are two issues. First, deceptive content needs to be truth-evaluable. It needs to be the kind of things that speakers can fail to believe, etc. So if *What Is Said* is so minimal that it fails to be truth-evaluable, there will be clear cases of lies that won't be classified as such. Second, they will *underclassify*. It will be very hard for an instance of deception to count as a lie.

ings of the utterance's constituent parts combined in accord with the utterances syntactic structure. It cannot be, then, that one *lies* when one utters something whose context-independent linguistic meaning is false (or, one believes to be false; or, one knows to lack good reasons to believe), and one *merely misleads* when the context-*dependent* content that's generated by making that utterance in that conversational context is deceptive. What indexical sentences --- and perhaps even more persuasively, *demonstrative* sentences --- show is that conversational context can be thought of as playing (at least) two different roles in communication: first, the conversational context acts upon what is uttered resulting in *What Is Said*; second, the conversational context interacts with *What Is Said* to ultimately determine *What Is Communicated*.<sup>6</sup>

Third, given that conversational context must, at least sometimes, play an important role in determining *What Is Said*, it is far from clear to what extent. Roughly, it is not obvious how to draw the line between *What Is Said* and *What Is Communicated* in a principled, non-arbitrary way. Moreover, it is not obvious how to draw line between the two in way that can make good on our intuitive judgments about what counts as a lie and what doesn't --- indexical sentences uttered in some contexts will be lies, uttered in others will only be merely misleading; and likewise for sentences containing demonstratives, completion cases, cases of expansion, etc. Let me sharpen the point. Consider the following *completion case*.

(4) Alan isn't ready.

---

<sup>6</sup>The conversational context might be thought to play a third role in determining *What Is Communicated*. I open my mouth, noises are produced. The conversational context might be thought to interact with these mouth-noises to determine *What Is Uttered*. Elliptical constructions are potentially examples of this. Here's an example.

(1) Alan is a wonderful person, and Beth is [a wonderful person] too.

The first half of (1) alters the conversational context in such a way that it can go on to determine the elided material in the second half. (The phenomenon is perhaps even more acutely brought-out by Discourse Ellipses). If this is right, then conversational context plays a role at the syntactic level (determining *What Is Uttered*), at the semantic level (determining *What Is Said*), and the pragmatic level (determining *What Is Communicated*).

Also, although it's been presented this way, don't think of these stages as actually happening in distinct chunks, one after the other. It very well may be that all three steps happen simultaneously. The distinction is meant to be conceptual, not temporal.



What does (4) say? Opinions vary.<sup>7</sup> This seems true. We could fill out the story surrounding an utterance of (4) according to which the utterance would, intuitively, be a lie. Let the story be one in which some activity  $\phi$  is clearly salient to all of the conversational participants. Furthermore, let the story be on according to which Alan, as a matter of fact, is very much ready to  $\phi$ . And I know this. My uttering of (4) is clearly a lie.

Now consider the following *case of expansion*.

(5) I went to the store and I got some eggs.

In many normal contexts, an utterance of (5) will (at the very least) *communicate* that I went to the store and got some eggs *from aforementioned store*. If I didn't get the eggs from the store --- if I, e.g., found them on the street (Hey, they looked perfectly fine! Don't judge me.) --- then is my utterance of (5) a *lie*? You might think: no.<sup>8</sup> I very well may have *misled* led to think that the eggs are fresh from the supermarket by *suggesting* as much --- but I didn't lie to you because I didn't *say* that. (Or, at least, so you might think).

The upshot is this. In order for the distinction between *What Is Said* and *What Is Communicated* to track our folk intuitions about when a deceptive utterance is a lie and when it is merely misleading, the contribution conversational context makes in determining *What Is Said* must (at least sometimes) allow for *completion* but not *expansion*.

---

<sup>7</sup>Here's a quick and incomplete survey. Those who are happy allowing extensive (although perhaps, constrained) contextual contributions to what is said by an utterance hold that what (4) says will depend on the context in which it is uttered. How so? Again, opinions vary. Very roughly, however, in order for an utterance of (4) to be appropriate, some activity  $\phi$  will be conversationally salient. If so, normally, (4) says *Alan isn't ready to  $\phi$* . Those who are weary of contextual contributions will say something else. Emma Borg, for example, holds that (4) says *there's something such that Alan isn't ready for it*. On the other hand, Cappelen & Lepore hold that (4) simply expresses the proposition that *Alan isn't ready*. What proposition is that, exactly? Don't ask me. I have no idea. (Exactly what would the world have to be like in order for this "proposition" to be true?) Finally, views according to which semantic content is *extremely skeletal*, like Bach's, will say that (4) doesn't express a complete proposition at all --- but rather a "proposition radical" in need of saturation from context to become something properly truth-evaluable.

<sup>8</sup>Admittedly, my intuitions start to become quite muddled. As will become important in a moment, however, I think we can fill in the surrounding story in different ways to evoke *both* intuitions.

## 2.3 Saul's Account of the Distinction

The lessons from the last section --- and, in particular, the observation that a wedge needs to be drawn between the kind of contextual contribution found in *completion cases* and the kind found in *expansion cases* --- leads Saul to offer the following account of *What Is Said*.<sup>9</sup>

(NTE) A putative contextual contribution to what it said is part of *What Is Said* only if without this contextually supplied material, [the utterance] would not have a truth-evaluable semantic content in *C*.

Saul's account requires *What Is Said* to be truth-evaluable (and thus the kind of thing that can be falsely believed, and believed to be false) while drawing a mostly-clear, principled, and non-arbitrary line in the proverbial sand between the extent to which conversational context contributes to *What Is Said*. Conversational context does just enough work to render the utterance truth-evaluable, and no more. Any additional work done by context belongs in the realm of pragmatics.

Because without *some* contextual supplementation sentences containing indexicals, sentences containing demonstratives, and completion cases will fail to express truth-evaluable propositions, the requisite contextual supplementation will contribute to *What Is Said*. This is good. Sentences containing indexicals, sentences containing demonstratives, and completion cases can, in the right circumstances, be lies. And because cases of expansion, like (5) do *not* require contextual supplementation to be truth-evaluable, the content context contributes is, on Saul's account, merely *suggested* and not *said*. This is good too. My devious utterance of (5) needn't be a lie.<sup>10</sup> Good, yes, but not great. I think there are a number of hard

---

<sup>9</sup>Saul 2012, ch. 3

<sup>10</sup>This is not the last word, of course. Saul rightly points out that her account leaves room for a range of unclear cases. More needs to be said about *how* the contextually supplied material is determined. Care must be paid in avoiding the telling of a story according to which too large of a role is given to what the speaker intends, and what the audience understands --- for the reasons discussed above. And there is a worry about cases in which context fails to determine a unique referent/completion&c. for an utterance that's, intuitively, a lie. Saul's solution is to complicate the account by adopting a *supervaluational* approach (and amending the relevant definitions accordingly). An utterance made

cases that tell against Saul's account. (Or are, at least, invitations to say more.) I will suggest that these hard cases suggest that we approach the Lying/Misleading Distinction, and the distinction between *What Is Said* and *What Is Communicated* is very different way

## 2.4 Hard Cases

Saul's account is good. But there are problems. (Then again, aren't there always?) There are cases in which Saul's account wrongly classifies *lies* as *mere misleadings*, cases in which her account wrongly classifies *mere misleadings* as *lies*, and case about which her account says I know not what.

### 2.4.1 Classifying Lies as Misleadings

There are cases in which in Saul's account seems to misclassify what are intuitively lies as mere misleadings. Here's one.

You and I are catching up over coffee. We haven't seen each other in a long time. You have fallen out of touch with our old circle of friends. I haven't. My relatively dull life, I worry, can't hold a candle to your adventurous one. I want you to think highly of me. And I know you always thought well of our old mutual friend, Alice.

(6) **You:** Whatever became of Alice?

**Me:** Alice and I got married and had children.

I *did* get married and have children, and so did Alice --- but not to, and not with, each other. I want you to think otherwise.

This strikes me as a lie. (Here's a way to pump the intuition. Were you to find out later that Alice and I were never married to each other, it would

---

in a context in which there fails to be a uniquely salient truth-evaluable content may still very well count as a lie just so long as, for each truth-evaluable content in the admissible range, the utterance is believed by the speaker to be false (or, known to be lacking the appropriate reasons to believe to be true). There still looms the difficult problem of saying something informative about how exactly context determines the admissible range of interpretations.

completely appropriate for you to think I lied to you). Because my utterance in (6) doesn't require any contextual supplementation in order to express a truth-evaluable proposition, on Saul's view, I've merely misled you. Maybe you agree with the verdict of Saul's account in this case. Here's another case (lifted from Stokke forthcoming).

Jasper's neighborhood recently put on a Community Week. People helped their neighbors out with various chores and tasks that needed doing. Selfishly, however, Jasper used Community Week to fix the roof on his own house, ignoring the neighbors. The following week Jasper is having dinner with Doris. Jasper is keen to give Doris a good impression of himself.

(7) **Doris:** So how did you help out during community week?

**Jasper:** I fixed a roof.

Jasper's reply strikes me as a lie. In order for Saul's view to get this right, there must be something that Jasper *says* which he believes to be false. But there isn't. It is true, and Jasper believes it to be true, *that Jasper fixed a roof*. The deceptive content --- that Jasper helped out during Community Week by fixing someone *else's* roof --- is not part of *What Is Said* on Saul's account.

#### 2.4.2 Classifying Misleadings as Lies

There are cases in which in Saul's account seems to misclassify what are intuitively lies as mere misleadings. Here's one (again, lifted from Stokke forthcoming).

Larry is keen on making himself seem attractive to Alice. He knows she's interested in logic --- a subject he himself knows very little about. From talking to her he has become aware that she is under the mistaken impression that he has just finished writing a book. Larry has indeed been walking around with a manuscript for a book about logic. And he knows Alice has seen him with it. However, it's not a manuscript for a book he wrote himself, but rather one that he has been assigned to design a cover for by the publisher he works for.

- (8) **Alice:** Do you know a lot about logic?  
**Larry:** My book is about logic.

Certainly Larry has uttered something misleading. But has he *lied*? It doesn't seem to me that he has. His utterance certainly *suggests* that the book in question is one he has written (and, consequently, that he indeed does know a lot about logic). What does Saul's account say about cases like this? It's unclear. Stokke argues that on Saul's view *What Is Said* by Larry's utterance is believed by him to be false --- and thus, contrary to our intuitions, counts as a lie.

It isn't clear to me that Stokke is right about this. A proponent of Saul's position might point out that the situation is one in which there fails to be a uniquely salient truth-evaluable proposition expressed by Larry's utterance. Here are two candidate propositions.

- (9) a. The book *Larry has written* is about logic.  
b. The book *Larry has been walking around with* is about logic.

Larry believes (9a) to be false, but he also believes (9b) to be *true*. Because 'my' is an indexical, Larry's utterance in (8) requires some contextual supplementation in order for what he says to be truth-evaluable. If the conversational context supplements Larry's utterance so it expresses (9a), then Larry has lied. If, on the other hand, the conversational context supplements Larry's utterance so it expresses (9b), then Larry hasn't lied. If both (9a) and (9b) are admissible interpretations of Larry's utterance, then Larry hasn't lied. Is (9b) among the range of possible truth-evaluable propositions? It seems so.

Stokke, however, finds this line of argument unpersuasive. He says the following.

If Larry's utterance is indeterminate across a range of acceptable completions, it is hard to see how his utterance could be misleading. Clearly, the reason Larry is being misleading is because he intends to make Alice believe [(9a)], and as a result [that he knows a lot about logic]. Indeed, Alice will take him to be conveying both. So ... [(9a)] is certainly a putative contribution to what is said. Hence, (NTE) would seem to predict that Larry *says* [(9a)].

[Saul's] view at the very least must find a way of explaining why [(9a)] is not *said*, given the Minimal truth-evaluability principle.

One immediate thing to say on Saul's behalf in response: the reason that (9a) is not *said* is because it does not have a claim on being the unique truth-evaluable proposition expressed by Larry's utterance. Larry's utterance is indeterminate across a range of acceptable completions. But why's that? (There's a worry that if our intuitions about lying / misleading are playing too large a role in determining when an utterance is indeterminate across a range of propositions, then we aren't really in the business of *explaining* the Lying / Misleading Distinction in terms of *What Is Said*. And if that's not what we're doing, what are we doing?)

### 2.4.3 Ambiguity Between Direct and Indirect Speech

Here's a case, like the previous one, where things seem muddled.

I want you to think that Alan doesn't like you. As it happens, Alan isn't too keen about me. In fact, he said the following to me.

(10) You are getting on everyone's nerves.

You know that Alan and I had a talked to each other. But you don't know what Alan said to me. We have the following exchange.

(11) **You:** What did Alan say to you earlier?

**Me:** Alan said you are getting on everyone's nerves.

Did I lie to you? One might think: no. My utterance is ambiguous --- and intentionally so.

- (12) a. Alan said [that] you are getting on everyone's nerves.  
b. Alan said "You are getting on everyone's nerves."

The former disambiguation is one on which I've offered an *indirect* speech report, while the latter is one on which I've offered a *direct* report. Because

we don't (always, and straightforwardly) need to pronounce the quotation marks, and because we can elide the 'that', utterances like mine in this example are ambiguous.

I have a moderately strong intuition that, although I am being supremely sneaky, I haven't lied to you. But it is unclear to me what Saul's account of the distinction would say about a case like this. On her account, remember, *What Is Said* by an utterance is determined by the most minimal contextual supplementation required to make the utterance truth-evaluable. But this is a case in which it is not clear which utterance was *uttered*.<sup>11</sup>

#### 2.4.4 Scalar Implicature

Here's yet another class of potentially problematic cases.

You baked everyone a batch of cookies. They go quickly. So quickly, in fact, that you don't get a chance to have any. This understandably makes you sad. You notice that I am covered in cookie crumbs, and have a bellyache.

- (13) **You:** Did you eat all the cookies?  
**Me:** I ate some of the cookies.

In fact, I ate *all* the cookies.

Did I lie to you? It seems to me that I did. But if the literal meaning of 'some' is *at least one* (and not *some but not all*), as is standardly thought, then, on Saul's account, I've only uttered something that's merely misleading. (No contextual supplementation is needed to render my utterance truth-evaluable). My suggestion that I ate some and not all of the cookies is determined by pragmatics, not semantics.

One might think that this just shows that, in fact, scalar implicatures have been misclassified as a pragmatic phenomenon --- and that we should take 'some' to literally mean not *at least one* but rather *some but not all*. This won't do, though, either. We can change some key details in the story, and the question you ask (to something like, "How was your snack?") so as to

---

<sup>11</sup>See fn 6.

elicit the intuition that my answer ("I ate some of the cookies") is a merely misleading --- and not a lie.

#### 2.4.5 Presuppositions

There's a lively, far-from-settled debate in linguistics about whether *presupposition* is a semantic or pragmatic phenomenon. Either way, there are cases in which uttering something that presupposes something one believes to be false (or, something one knows one lacks good grounds to believe to be true) will seem like lies, and cases in which doing so will not.<sup>12</sup>

### 3 Maintaining Plausible Deniability About Having Done Something Deceptive

In the previous section we looked at one way of tying the Lying/Misleading Distinction to the distinction between *What Is Said* and *What Is Communicated*: namely, the view Saul develops in Saul 2012. I hope, also, to convince you that there are reasons to worry about her account. At the very least, I hope to have convinced you that there's no harm in trying something else. Allow me to give it a go.

Here's what I think. The Lying/Misleading Distinction is roughly tracking what the speaker can, and cannot, maintain plausible deniability about having communicated. Roughly, an utterance is considered a *lie* when we think, were the fact that the speaker communicated something s/he believed to be false (or, knowing lacked good grounds to believe to be true) to be *found out*, s/he would not be able to maintain plausible deniability about having done something deceptive. There's no plausible story the speaker could spin according to which s/he was being as conversationally cooperative as s/he's able and was *not* being deliberately misleading.

On the other hand, an utterance is considered to be *merely misleading* when we think, were the fact that the speaker communicated something s/he believed to be false (or, etc.) to be found out, s/he *would* be able to plausibly

---

<sup>12</sup>Generating such cases is an exercise hereby left to the reader.



claim that she wasn't intentionally being deceptive --- even if as a matter of fact that isn't at all the case. Mere misleadings, then, occur in situations where the misleader could, after the fact, make a plausible case that what s/he communicated was *accidental* --- it was the product of a perfectly understandable (and perhaps even, faultless) misunderstanding.

***What Is Said:*** The communicated content about which the speaker cannot reasonably maintain plausible deniability about having intended to communicate.

Here's a way to think about it. Suppose the misled party later on discovers that the speaker has misled them. She confronts the speaker about it. Would the speaker be able to make a reasonable case act of misleading happened in good faith, was unintentional, and the unfortunate result of an understandable mistake. To wit, would the speaker be able to get away with saying something like "Oh! We misunderstood each other!" --- even though, of course, it would be false. The degree to which the speaker could reasonably get away with saying something like this is the degree to which the deceptive content is merely misleading as opposed to a lie.

Two quick worries. First, plausible deniability admits of degrees. Some stories are more plausible than others. So one might object to the correspondence between the distinctions on the following grounds. Lying does not admit of degrees; you either lie or you don't. Response? False. Lying admits of degrees too. Second, you might worry that my account gets things back to front. It is the fact that you *lied* that explains why you would not be able to maintain plausible deniability about having done something deceptive --- not the other way around! Response? Rather than address the worry head-on, allow me to try to persuade you indirectly by illustrating how the account address the various cases --- and in particular, the hard ones --- we've encountered so far.

### 3.1 Hard Cases Made Easier

In order to fully understand the proposal, we should apply it. We've looked at a number of cases. Consider the cases of expansion and completion. On Saul's account, if the deceptive content of an utterance is a result of *expansion*, then the utterance is merely misleading; if the deceptive content of an utterance is a result of *completion*, then the utterance is a lie. On this account, however, the conversational context can be such so that the deceptive content resulting from expansion can be a lie; and the conversational context can be such so that the deceptive content resulting from completion can be misleading. This, I contend, lines up with our intuitions about the cases in §2.4.

Let me attempt to warm you up to the idea by presenting you with a series of cases --- ranging from clear cases of lies, to murkier cases, to clear cases of mere misleadings. Consider the following utterance.

(14) There is absolutely no poison in your cup.

I utter (14) even though, as I know, your cup is chock full of poison. Did I lie to you? Obviously yes. And notice, were my deception to be unearthed, I would not have a leg to stand on. I would not be able to maintain any plausible deniability about having done something deceptive. To bring this out, suppose you were to discover that I knew your cup was full of poison.

(15) Oh! You misunderstood me. By 'absolutely no poison' I meant a whole bunch of poison.

This, obviously, wouldn't help me at all. The uttering of (15) is not able to "take back" the fact that what I said by uttering (14) was a *lie*. (Your post-hoc stipulation cannot all by their lonesome make words mean what you now want them to). It was reasonable to expect me to *know* that my words will communicate what they, in fact, did communicate to you (namely, that there's no poison in your cup). And *all it takes* for this expectation to be reasonable is that you take me to be a competent speaker of English. (And, one might think, my uttering of a sentence in English gives you a tremendous amount of evidence that I am a competent speaker of English).

**Uttering (14) is a Lie.** If I know that your cup is full of poison, then by uttering (14) either (i) I am being deceptive or (ii) I have made some kind of communicative mistake. Because, if I am a competent speaker of the language, it is not plausible that I have made some kind of communicative mistake in this case, it is reasonable to conclude that I was being deceptive.

This is a clear case of lying. It is also clearly a case in which the speaker cannot reasonably hope to maintain plausible deniability about having communicated something deceptive.

What about utterances containing Indexical and /or Demonstratives? Consider the following.

(16) I've never seen him before in my life!

You cannot reasonably hope to use (16) deceptively and avoid *lying* by, after the fact, claiming your utterance of (16) meant something like:

(17) Bob Dole has never seen him before.

Why? Because the rule associated with indexicals like "I" --- it's *kaplanian character* --- is *strict*. The referent of "I" in an utterance *u* made in context *C* is the speaker of *u* in *C* --- and that's that; there's no wiggle room. Moreover, insofar as you are a competent speaker of English, you are reasonably expected to *know* as much. Of course, not *all* indexicals (and even more so for demonstratives) are governed by a rule as strict as the one associated with the first-personal pronoun. In fact, you *could* reasonably hope to communicate something false with (16) and avoid lying by, after the fact, claiming that by uttering (16) you meant something like the following.

(18) I've never seen [someone other than whom I, sneakily, intended you to think of] before in my life!

How well this would work, of course, depends on how plausible it is that the conversational context in which (16) was uttered was one in which I

could reasonably have taken this other person to have been salient enough to have been the referent of "him". Context matters. If the conversational context was one about which it could have been reasonable for someone who is in the position the audience takes the speaker to have been in to have, in good faith, believed someone else was salient enough to be the referent of "him", then the speaker's utterance of (16) wasn't a lie. On the other hand, if the conversational context was one about which it just would *not* have been reasonable to think that someone else was salient enough to be the referent of "him" (a context, for example, in which it is super clear that we've been talking about *Alan*), then the speaker's utterance of (16) would be a lie.

**Uttering (16) Can be Either.** Whether an utterance of (16) is a lie, or merely misleading depends on features of the conversational context --- in particular, whether or not the context was such that it is plausible the speaker could have been mistaken about who would be the salient referent of "him".

The meaning of a lot of our words --- and consequently the meaning of a lot of our utterances --- depend crucially on facts about what is, and what is not, salient in the context in which the conversation takes place. And it's not implausible that totally reasonable, mostly faultless *misunderstandings* can result from differences in what the conversational participants take to be salient in the conversational context. Skillful misleaders are able to exploit this fact. Even though, in fact, there is no such misunderstanding, the skillful misleader is able to use the fact that there plausibly *could be* to communicate something s/he believes to be false (or, knows to be lacking good reasons to believe to be true) all while deviously maintaining plausible deniability about having done so.

One consequence of this way of thinking about the Lying/Misleading Distinction, as mentioned above, is that the difference between cases of *completion* and cases of *expansion* --- which plays a central role in Saul's account --- is somewhat of a red herring. Cases of completion can be either lies or misleadings --- depending on the features of context highlighted above. Consider the following completion case.

(19) Frank has had enough.

An utterance of (19) will be a lie when the conversational context is such that there is a salient way to complete it which is so obvious that it would be totally implausible for the utterer to be mistaken about it. On the other hand, if the utterer could make a plausible case that there was no such obviously salient completion, utterances like (19) in such contexts will be merely misleading. Likewise, cases of expansion can be either lies or misleadings --- again, depending on what features have been made salient, in an overwhelmingly obvious way, in the conversational context. Consider the following conversation involving a case of expansion.

I know that you hate when I bring home eggs that I've found on the street. So much of the food in our house has come from my scavenging that you, understandably annoyed with me, worry that we will have nothing respectable to eat for brunch.

(20) **You:** What are we going to eat for brunch? I refuse to eat any of the things you've found on the street.

**Me:** I went to the store and I got some eggs.

Did I lie to you? I think so. Why? The conversational context is such that there is no way I could maintain plausible deniability about communicating something highly misleading to you. I can imagine (because I've seen it oh so many times before) the incredulous-and-angry look you'd give me were I to respond to your realization that the eggs I used to cook your omelet had been found on the street by saying "I never *said* I got the eggs *from the store*. I was just reporting to you facts about what I did this morning." There is no plausible story according to which I utter what I did in (20), I am being conversationally cooperative, and I am genuinely not attempting to deceive you. Therefore, on my account, I've lied. And that seems right.

### 3.2 The Role of Context in Misleading

Just offhand, it seems as though it is, in general, easier to *merely mislead* (as opposed to *lie*) the greater the degree to which the conversational context

contributes to suggesting the misleading content. (This is not, of course, *exceptionless*. But it does seem to be roughly true).

How does this observation sit with the account developed above? It sits quite well. First, not always but usually, the greater the degree conversational context is operative in determining what is communicated by an utterance, the greater degree to which it will be possible that there are (reasonable) differences in what the conversational participants take to be common ground, conversationally salient, etc. If we imagine *explicitly* running through the steps it would take to calculate what is communicated by a particular utterance in a conversational context, the larger the role context plays, the more steps are needed, more assumptions are granted, etc. The thought is that the more conversational context needs to contribute in determining what an utterance communicates, the greater degree to which *What Is Communicated* by that utterance is epistemically fragile.

Here's another way in which heavy context-dependence correlates with misleading. It's harder, as a matter of psychological fact, to remember what all the features of the conversational context set were like. Even if it is clear to me *at the time* that *X* is salient and *p* is part of the common ground, etc., later on, as my memory fades, it very well may become less clear to me whether or not *X really* was as salient as I thought it was, etc. (Notice, for example, the role that *tone* plays in conversation --- and how difficult it is, after the fact, to remain confident in your convictions about what using that tone in that context was communicating). So, the more context matters, the easier it is for the misleading party to plausibly make the case that "this has all just been one big, but completely reasonable, misunderstanding!" Conversely, the *less* context matters in determining the deceptive content communicated by the misleading party, the harder it will be to maintain that what was communicated by the speaker was merely the product of a reasonable and faultless misunderstanding.

## 4 Ramification for Ethics: is lying morally worse than misleading?

Is it morally worse to lie than merely mislead? At first blush, the answer definitely seems to be yes.<sup>13</sup> But if, as I'm claiming, the Lying/Misleading Distinction just tracks the extent to which we can, or cannot, reasonably hope to maintain plausible deniability about having been deceptive, it is extremely hard to see why lying would be morally worse than being deliberately misleading. I contend that it is extremely hard to see why this could be because it isn't --- all else equal, lying isn't any morally worse than being deliberately misleading. We believe lying is worse, but we're wrong. What accounts for the fact, though, that we (apparently) all got this one so wrong? If it really is *false* that lying is worse than misleading, why is it that this erroneous belief is so deeply held, and widespread?

Settling aside my account of the distinction, there is considerable independent reason to think that lying isn't any worse than misleading. First, there are cases (like the Peanut example, borrowed from Saul 2012, reprinted at the beginning of this paper) in which being misleading seems no better than lying. Second, as Saul demonstrates by extensively surveying the most promising accounts of why lying is *ceteris paribus* worse than misleading, it is incredibly difficult to give a convincing principled account of the ethical distinction. When we reflect on the distinction, it doesn't seem to hold up to scrutiny. And if I'm right about how to think about the Lying/Misleading Distinction, not only is lying not morally any worse than merely misleading --- it actually looks like *misleading is morally worse than lying*. There is something very ugly about, on top of being deceptive, sneakily engineering situations so as to evade the detection of wrongdoing. The kind of skill required to be an expert-level misleader is almost *sociopathic*. To drive the point home, compare the situation of the victim of misleading to the situation of the victim of lying --- were they to discover that the deceiver had communicated something the deceiver lacked good reasons to believe. When the deceptive content is a lie and the victim makes such a discovery, given my account of the distinction, the victim now *knows* that the deceiver has been deceptive. The victim knows that

---

<sup>13</sup>Keep in mind that the question is: is it morally worse to deliberately lie than it is to be *deliberately* misleading.

s/he has been wronged by the deceiver. On the other hand, when the deceptive content is merely misleading and the victim makes such a discovery, the deceiver is able to *conceal* the fact to the victim that s/he has been wronged. In fact, depending on his/her manipulative skills, the deceiver may be able to push at least some of the blame for the "misunderstanding" onto the victim.<sup>14</sup> There's something extra bad about this --- the victim of the misleading is being wronged twice over.

If lying isn't any worse than misleading, what could possibly explain our belief to the contrary?

#### 4.1 Why Do We Prefer to Mislead Rather Than Lie?

I am going to argue that, although it isn't in general morally any better to mislead than it is to lie, it is completely reasonable to expect creatures like us --- profoundly social beings whose success crucially depends on our ability to coordinate with each other --- to come to internalize a moral preference for misleadingly over lying.

The story turns on an observation of Kant's (filtered through David Lewis). Kant famously argued that we have a duty not to lie --- in part because if everyone were to lie whenever they felt like (i.e., if the maxim were to be universalized), communication itself would irrevocably breakdown rendering it impossible to lie, which presupposes that people, generally, believe what we tell them. David Lewis claimed that in order for there to be a signaling practice rich and robust enough to be a *language* there must be a convention of *truthfulness* and *trust* among the language's members. Roughly, this means that in order for a population to form a linguistic community, the members of the population have to more-or-less speak truly, and trust that other's are speaking truly as well. If Lewis (and Kant) are correct, a community of habitual liars wouldn't be able to establish a signaling convention rich and robust enough to be a language (and, thus, wouldn't be "a community of liars" after all).

---

<sup>14</sup>This would be a situation that shared many of the important features with the phenomenon of *Gaslighting*. In some cases of skillful misleading, the victim is fairly sure of what the deceiver is communicating *during the conversation*. But later, when the victim confronts the deceiver, the victim's *true* belief about what was communicated is undermined by the deceiver's claim that there was a misunderstanding.



Because (i) communities that are able to successfully coordinate with one another will, all else equal, do better than communities that aren't and because (ii) having a rich and robust signaling convention facilitates successful coordination, communities lacking a critical mass of liars will fair much better than communities up to their ears in them. We can expect, therefore, that via a process of Social Evolution, the successful communities would come to be comprise members who've internalized a strong non-instrumental desire to not lie.

But wait. Even if this story is on the right track, all that this shows is that it's not unreasonable to expect creatures like us (who are the products of eons of Social Evolution) to prefer lying to *not lying*. This doesn't explain how or why we would come to internalize a preference for lying over *misleading*. In fact, on might think, wouldn't *the very same story* predict that social creatures would come to internalize a strong non-instrumental desire to not mislead, as well?

No. Not if my account of the Lying/Misleading Distinction is correct. Here's why. Lying, if found out, undermines the convention of truthfulness and trust necessary for the health of a successful language. This isn't the case for merely misleadings. Because, on my account, if you merely mislead, you thereby maintain plausible deniability about having done something that would undermine the truthfulness and trust which is required for a flourishing language.

If all this is correct, however, it leaves us in a bit of a bind. On the one hand, it doesn't seem like we have any good *reason* to prefer misleading to lying. (In fact, maybe quite the opposite). On the other hand, it is good, conditional on everyone else having it as well, that we have such a preference. What are we to do. I'd recommend trying to avoid being deceptive in the first place.