

Lying, Misleading and Fairness¹

Emanuel Viebahn, Humboldt-Universität zu Berlin

emanuel.viebahn@hu-berlin.de

Forthcoming in *Ethics*

Abstract

Sam Berstler defends a general moral advantage for misleading over lying by arguing that liars, but not misleaders, act unfairly towards the other members of their linguistic community. This article spells out three difficulties for Berstler's account. Firstly, though Berstler aims to avoid an error theory, it is dubitable that her account fits with intuitions on the matter. Secondly, there are some lies that do not exhibit the unfairness Berstler identifies. And, thirdly, fairness is not the only morally relevant difference between lying and misleading.

1. Introduction

Is it morally better to mislead than to lie? And, if so, why? In “What’s the Good of Language? On the Moral Distinction between Lying and Misleading,” Sam Berstler answers these questions by presenting an intriguing new account of the intuitive moral difference between lying and misleading.² Berstler's account consists of two main claims. First, she argues that liars, but not misleaders, misuse the conventions of truthfulness and trust that fix the meanings of our language. For this reason, liars act unfairly towards the other members of their linguistic community in a way that misleaders do not. Then, Berstler tries to make plausible that this unfairness is the *only* morally relevant difference between lying and misleading. Taken together, these claims imply that it is always worse to lie than to mislead (holding all else fixed) – a result which, according to Berstler, matches our intuitions on the matter.

¹ For very helpful comments and discussion, I would like to thank Sam Berstler, Marius Drozdowski, Felix Timmermann, Derya Yürüyen and two anonymous associate editors for *Ethics*.

² Sam Berstler, “What’s the Good of Language? On the Moral Distinction between Lying and Misleading,” *Ethics* 130 (October 2019): 5–31. Henceforth *WGL*.

In what follows, I will point to three difficulties for Berstler's account. To begin with, I will mention two uncertainties concerning the intuitions Berstler aims to capture. Then, I will argue that fairness-based considerations cannot support a general moral advantage for misleading over lying, as it is possible to lie while speaking non-literally, and non-literal lies do not exhibit the unfairness Berstler identifies. And, finally, I will try to show that even if fairness is one factor that can set apart lying and misleading, it is not the only morally relevant difference between the two options: lying and misleading can also differ in terms of commitment, and this difference can be morally relevant, too. Even if we assume that lying is unfair in a way that misleading is not, we should thus resist the claim that it is always worse to lie than to mislead.

2. Berstler's account of the moral difference between lying and misleading

Before turning to my critical points, I would like to introduce the core claims of Berstler's account: the cases it is meant to apply to, the intuitions about these cases it aims to capture, and the way in which it is meant to capture these intuitions.

Berstler is interested in cases in which someone faces a choice between lying and misleading. But how must the notions of *lying* and *misleading* be spelled out here? Berstler accepts the traditional assertion-based definition of lying, according to which to lie is to assert something one believes to be false (WGL: 8). And she clarifies that to mislead, in the sense relevant here, is to intentionally conversationally implicate something one believes to be false by way of asserting something one does not believe to be false (WGL: 7).³ Furthermore, the only difference between the options of lying and misleading is meant to be the way in which a certain believed-false proposition is put forward: whether it is asserted or conversationally implicated. Everything else about the situation is held fixed, including *which* believed-false proposition is put forward. So Berstler is interested in situations in which there is a choice of either asserting a proposition one believes to be false or conversationally implicating the same believed-false proposition through an assertion that is not believed to be false.

Berstler discusses several examples of this kind, and it will be helpful to introduce one of them here:

³ To be fully precise, Berstler is only interested in cases involving *particularised* conversational implicatures, which I will henceforth simply refer to as *conversational implicatures*. Berstler does not explicitly mention that the conversational implicature has to be generated through an assertion that is not believed to be false, but these are clearly the kinds of cases that are at issue.

Judith has just taken a bite of her brother Quentin’s cake. Quentin asks her, “What do you think?” Painfully sensitive to hurting others’ feelings, Judith can’t bring herself to tell Quentin what she thinks—namely, that the cake is disgusting. But Judith also finds that she cannot force the words “It’s delicious!” from her mouth. So Judith splits the difference. “This is the best dessert I’ve had all week,” she says. In fact, Judith has eaten only one other dessert this week: a horribly burned, failed baking experiment. Quentin’s cake still tasted better than that. So, what Judith has said is literally true. (WGL: 1)

Judith is faced with a choice between uttering (1) and (2):

- (1) It’s delicious.
- (2) This is the best dessert I’ve had all week.

The believed-false proposition Judith puts forward in this case is that the cake is delicious. And the choice Judith faces is between asserting this proposition through uttering (1) or conversationally implicating the same proposition by uttering (2), which she believes to be true.

Which intuitions about such choices does Berstler aim to capture? Berstler claims that in the case introduced above and “[g]enerally, we intuit that to lie is worse to mislead” (WGL: 6). She does, however, point to some high-stakes cases in which the intuition disappears (WGL: 26). As an example of such a situation, she mentions Saul’s case of the attempted peanut murder, where it seems to make no moral difference whether the speaker lies or misleads in order to get a person with a fatal peanut allergy to eat a meal containing peanut oil.⁴ In high-stakes situations of this kind, Berstler holds, the light moral difference between lying and misleading is swamped by the heavy wrong of an attempted murder. However, throughout the paper it is clear that Berstler aims to capture a *general* preference for misleading over lying that is only absent in certain high-stakes cases.

⁴ Cf. Jennifer Saul, *Lying, Misleading, & What is Said* (Oxford: Oxford University Press, 2012), at pp. 72–73. Berstler also discusses a situation involving lower stakes, in which an agent faces a choice between lying or misleading about a friend’s hairdo, which she hates. Here, Berstler argues, an intuitive preference for misleading might be swamped because “both the act of lying and the act of misleading [appear] so obviously permissible” (WGL: 27). I am not sure Berstler should adopt this explanation, as almost all cases discussed in the literature feature a choice between two utterances that are obviously permissible or two utterances that are obviously impermissible. So, to accept that in such cases there is no intuitive preference for misleading (possibly because it is swamped by the obvious permissibility/impermissibility) is close to accepting that there is no general intuitive preference for misleading, which would go against Berstler’s overall outlook.

In order to account for this presumed intuition, Berstler introduces the Lewisian account of how the meanings of a language are connected to conventions governing the practice of assertion.⁵ According to Lewis’s metasemantics, assertion is subject to the conventions of truthfulness and trust, which Berstler puts as follows (WGL: 18):

Let x be a sentence and p a proposition. If x (conventionally) means p , then:

Truthfulness: (a) Speakers assert x only if they believe p .

Trust: (b) If speakers assert x , their addressees come to believe p (unless they have some other reason not to).

It is because of these conventions that the expressions of a language have the meanings they have. And it is because the expressions have their meanings that the members of a linguistic community can use language to transmit information and coordinate. The members of a linguistic community thus benefit from the conventions being in place, i.e. from the fact that asserters usually conform to the conventions of truthfulness and trust.

Berstler also shows that there are no corresponding conventions of truthfulness and trust that link the meanings of a language to the propositions conversationally implicated. Implicatures (of the particularised conversational kind that is at interest here) are not associated in a regular way with certain sentences or meanings, but have to be calculated on an *ad hoc* basis. Conversational implicatures are thus not governed by the kinds of meaning-constituting conventions that apply to assertion (WGL: 18–20).

Next, Berstler argues that this metasemantic difference between assertion and conversational implicature is morally relevant. In particular, she argues that assertion as a mutually beneficial practice obligates the members of a linguistic community to conform to the conventions of truthfulness and trust. In disregarding these conventions, liars thus act unfairly towards the other members of their linguistic community. Misleaders, by contrast, do not act unfairly in this way, as they do not misuse conventions that are of metasemantic importance.⁶

Finally, Berstler tries to show that the unfairness of assertion is the *only* morally relevant difference between lying and misleading. On the one hand, she attacks several existing accounts of such

⁵ David Lewis, *Convention: A Philosophical Study* (1969; repr., Cambridge, MA: Harvard University Press 2002).

⁶ See WGL: 20–25. Note that Berstler’s view is not that liars *damage* or *threaten* the beneficial practice of assertion through misusing the conventions of truthfulness and trust. Rather, the wrong committed by liars (but not misleaders) is an act of unfairness. Theorists have argued for related views, on which liars do damage the practice of assertion, cf. Alan Strudler, “The Distinctive Wrong in Lying,” *Ethical Theory and Moral Practice* 13 (2010): 171–79, and Seana Shiffrin, *Speech Matters: On Lying, Morality, and the Law* (Princeton, NJ: Princeton University Press, 2014).

a difference, arguing either against the difference posited or against its moral relevance (WGL: 11–17). On the other hand, she highlights the similarities of assertion and conversational implicature, which in her view “have similar epistemic profiles and play identical communicative roles” (WGL: 25). This leads Berstler to conclude that “the only salient difference is that assertion depends on conventional language, whereas conversational implicature also depends on certain kinds of ad hoc reasoning” (WGL: 25–26). When faced with a choice between lying and misleading, speakers thus “always have one extra reason not to lie” (WGL: 7) – all things considered and holding everything else fixed, lying is always worse than misleading.⁷

In the next three sections, I will pinpoint three problems for Berstler’s view. While Berstler may well be right to highlight linguistic unfairness as one kind of wrong that is associated with many cases of lying but not with misleading utterances, I am sceptical that this account can explain the intuitive moral difference between lying and misleading.

3. Intuitions about lying and misleading

To begin with, I would like to discuss two uncertainties about the intuitive basis that Berstler’s account aims to capture, which could lead to two ways in which the account diverges from intuitions on the matter.

Firstly, there is reason to doubt that we do in general intuit that lying is worse than misleading. Of course, there are the high-stakes cases Berstler mentions, in which a small intuitive preference for misleading is plausibly swamped. But even beyond such high-stakes cases, there are cases for which theorists have noted that lying is not intuitively worse than misleading. To begin with, there are cases in which the choice between lying and misleading does not seem to matter. For example, Holger Baumann describes a case in which someone chooses between a lie and a misleading utterance in order to keep her partner in the dark about an affair, and notes that in this case the speaker

⁷ Berstler’s arguments for the similarity of assertion and implicature can be understood as arguments for a certain view of what must be held fixed in morally evaluating cases of lying and of misleading. Of course, all parties to the debate accept that lies should be compared to their direct counterparts of misleading – we are only allowing differences that result from the agent’s choice between lying and misleading. But as there are different views of the nature of lying and misleading (and their relation), there are also different views of what is held fixed in the relevant comparisons. By highlighting the communicative similarities of asserting and implicating, Berstler tries to make plausible that in comparing lying and misleading, everything must be held fixed that does not have to do with metasemantic conventions. Below, I will argue that such a view of what must be held fixed is too strict. We must allow for other differences that follow from differences between lying and misleading, in particular for differences in communicative commitment.

is not doing anything morally better by opting for the misleading utterance.⁸ In a similar vein, Saul discusses a variant of the well-known case in which a dying woman is kept in the dark about the recent death of her son. In the variant, the person choosing to lie to or mislead the dying woman does so only to inherit the dying woman's fortune. Saul holds that it is unlikely that misleading is thought to be better than lying in this variant of the case.⁹

Furthermore, there are cases in which misleading is intuitively *worse* than lying. As an example of such a case, Felix Timmermann and I have described the following situation involving a willingly deceived addressee:

John is an ambitious and passionate amateur chef, but unfortunately not blessed with too much talent. His wife, Joanne, has repeatedly had the experience that John got quite upset when she showed too little enthusiasm about his culinary achievements, or even criticized some meal he prepared. From this she has drawn the conclusion that John prefers an insincere compliment to an honest appraisal of his products. When John once again puts too much salt in the soup and asks whether she enjoyed the meal, Joanne could either utter the false [3] or the true but misleading [4]:

[3] It was the best fish soup I've ever tasted. It was delicious!

[4] My mother used to cook that soup on Christmas Eve. I loved it.¹⁰

Here it seems that an attempt to mislead by uttering (4) is worse than the lie (3), possibly because the attempt to mislead could make John suspicious and uncover the deception.¹¹ These are neither high-stakes cases, nor are they exceptional. Still, there appears to be no intuitive advantage for misleading.

A second uncertainty about the intuitions Berstler aims to capture concerns cases in which we *do* think that lying is worse than misleading, such as the case of Judith and Quentin. In considering such cases, we arguably not only have the intuition *that* lying is worse, but also intuitions about *why* lying is worse. And here it seems to me (and many of those I have spoken to) that the difference between lying and misleading has to do with how the speaker acts *towards the addressee*. In particular, one disadvantage of lying seems to be that the liar wrongs *the addressee* in a way the misleader avoids.

⁸ Holger Baumann, "Gibt es einen moralisch relevanten Unterschied zwischen Lügen und Irreführen?" *Zeitschrift für praktische Philosophie* 2 (2015): 9–36, at p. 25.

⁹ Saul, *Lying, Misleading, & What is Said*, at p. 87.

¹⁰ Felix Timmermann and Emanuel Viebahn, "To lie or to mislead?" *Philosophical Studies* (2021): 1481–1501, at p. 1486. Saul, *Lying, Misleading, & What is Said*, at p. 90, describes a real-life case in a political context in which misleading appears to be worse than lying.

¹¹ In discussing cases of this kind with others, I have found that a majority sees a moral advantage for misleading, though I have also encountered reactions that deny such an advantage.

Berstler's account goes against this intuition: it entails that the liar wrongs *the members of the linguistic community* in a way the misleader avoids. Of course, in doing so the liar also wrongs the addressee as a member of the linguistic community, but only indirectly and to the same extent as every other member of that community. But there is no clear intuition (at least in my experience) that liars wrong uninvolved people, and that they wrong uninvolved people to the same extent as their addressees. By contrast, I have repeatedly encountered the intuition that the choice between lying and misleading matters *more* for the addressees than it does for uninvolved people.

Of course, some recalcitrant intuitions are not a reason to rule out an account. But I do think that intuitions do not pan out quite as neatly as Berstler assumes (and as is often assumed in the debate): if a broad range of examples is considered, a general preference for misleading over lying is doubtful even once high-stakes cases are bracketed. Furthermore, even in cases in which there is an intuitive preference for misleading, these intuitions may not fit perfectly with Berstler's account of *why* misleading is better. This suggests that Berstler, who hopes to avoid an error theory about folk intuitions (WGL: 2), might be arguing for a view that conflicts with intuitions after all. It also shows that there is a need for empirical studies about folk intuitions to set the debate on stabler foundations.¹²

Let us set these uncertainties about the intuitions aside for now and turn to Berstler's view that all lies are associated with a specific kind of metasemantic unfairness, which is not to be found with misleading utterances.

4. Lying beyond conventions

On Berstler's view, liars act unfairly because they misuse metasemantically relevant conventions governing the practice of assertion. Misleaders do not act unfairly in this way because these conventions do not govern the practice of conversationally implicating. If this difference in unfairness is meant to be a *general* difference between lying and misleading, as Berstler argues, then the metasemantic unfairness has to be present in all cases of lying and no cases of misleading. The lying-misleading distinction has to align with the distinction between the conventional communicative

¹² For an initial study in this area, see Alex Wiegmann and Neele Engelmann, "Is lying morally different from misleading? An empirical investigation," in *From lying to perjury: Linguistic and legal perspectives on lies and other falsehoods*, ed. Laurence R. Horn, (Berlin: De Gruyter Mouton, forthcoming).

practice of assertion and the non-conventional communicative practice of conversationally implicating. Berstler indeed accepts such a view of the lying-misleading distinction:

To lie is to make an assertion that one believes not to be true. [...] When I assert p, I say p in order to mean p, and I intend to contribute p to the conversation. [...] [I]o mislead is to make a certain kind of implicature that one does not believe to be true. (WGL: 8–9)

But there are reasons to believe that there is no such alignment, as I will now argue: conversational implicatures can be used to assert and to lie, and so liars need not misuse metasemantically relevant conventions of assertion. Some lies are not unfair in the way Berstler argues.

The structure of the argument can be put as follows. The first premise emphasises that the Lewisian conventions of truthfulness and trust apply only to (conventional) assertion, and not to conversational implicature. In the second premise, it is argued that it is possible to lie non-literally, i.e. with conversational implicatures. This leads to the conclusion that the metasemantic unfairness Berstler associates with lying is not present in all cases of lying, and thus cannot ground a general moral difference between lying and misleading.

The key point of the first premise is that the Lewisian metasemantic conventions apply only to *conventional* assertion: they apply to cases in which speakers assert p by uttering a sentence that conventionally means p. The standard case of conventional assertion is one in which a speaker asserts p by uttering a sentence that has p as its semantic content. But Berstler (WGL: 31) holds that *conventional* implicatures can also be used to conventionally assert, and that seems right. In such a case, a speaker asserts p by uttering a sentence that conventionally implicates p. Importantly, however, conventions of this kind do not apply to conversational implicature – the practice of conversationally implicating is not metasemantically relevant. Berstler nicely sums up why:

Why doesn't [conversational] implicature play a metasemantic role? The answer lies in the fact that the relevant kinds of implicatures, Gricean particularized conversational implicatures, just aren't cases of conventional meaning. Since they aren't cases of conventional meaning, we don't use them with any regularity, and, a fortiori, these regularities cannot explain their meaning. (WGL: 19)

So, the unfairness Berstler associates with lying is found with conventional assertion, but not with conversational implicature.

The second premise highlights the possibility of non-literal lies to call into question the assumed alignment between the lying-misleading distinction and the distinction between conventional

assertion and conversational implicature. If we take a brief look at recent accounts of the theoretical difference between lying and misleading, the alignment-assumption seems to be borne out. Don Fallis, Jennifer Saul, Andreas Stokke and others hold that in order to tell a lie with content *p*, a speaker has to utter a sentence that has *p* as its semantic content.¹³ If someone insincerely puts forward *p*, but does so without uttering a sentence that has *p* as its semantic content, but rather through conversationally implicating *p*, that is a case of misleading.

However, there are reasons to doubt this view of the lying-misleading distinction. On the one hand, it has been argued that assertion need not be literal: that it is possible to assert *p* through conversationally implicating that *p*. For example, Merrie Bergmann argues for the possibility of *metaphorical* assertion by discussing utterances of the following kind:

- (5) The nuclear reactor is a tinderbox.¹⁴

Bergmann holds that a speaker can use (5) to assert that the nuclear reactor is likely to fail, although that is not the semantic content of the sentence uttered, but rather a conversational implicature (according to a Gricean account of metaphor). One reason to think that this is the right verdict can be brought out by considering the commitment a speaker incurs by uttering (5). Following Peirce and Brandom, assertion is commonly tied to commitment: in asserting *p*, one commits oneself to *p* and thus takes on a responsibility to justify (or defend) *p* if challenged.¹⁵ In line with this, the speaker of (5) clearly commits herself to the reactor being likely to fail, taking on a responsibility to justify that content. Now, if there can be metaphorical and thus non-literal assertion, that gives us a reason to accept the possibility of non-literal *lying*. After all, it is widely held (and accepted by Berstler) that lies are believed-false assertions.¹⁶

¹³ Don Fallis, "What is Lying?" *The Journal of Philosophy* 106 (2009): 29–56; Saul, *Lying, Misleading, & What is Said*; Andreas Stokke, *Lying & Insincerity* (Oxford: Oxford University Press, 2018). Berstler does not put her view of lying in terms of semantic content, but she holds that telling a lie with content *p* requires speakers to utter a sentence that conventionally means that *p* (WGL: 17–19). As noted above, this includes cases in which the content of the lie is the semantic content of the sentence uttered and cases in which it is conventionally implicated. The possibility of lying with conventional implicatures is discussed and accepted by Andreas Stokke, "Conventional Implicature, Presupposition, and Lying," *Aristotelian Society Supplementary Volume* 91 (2017): 127–147.

¹⁴ Merrie Bergmann, "Metaphorical assertions," *The Philosophical Review* 41 (1982): 229–245, at p. 231.

¹⁵ Charles Sanders Peirce, "Judgment and assertion," in *Collected papers of Charles Sanders Peirce Vol. V* (Boston: Harvard University Press, 1934), 385–387; Robert Brandom, "Asserting," *Notis* 17, 4 (November 1983), 637–650.

¹⁶ Bernard Williams, *Truth and truthfulness: An essay in genealogy*, (Princeton: Princeton University Press, 2002), at pp. 98–99, also argues in favour of metaphorical assertion, although elsewhere he seems to accept the alignment-assumption.

On the other hand, theorists have directly argued for the possibility of non-literal lies. A first example of a metaphorical lie is discussed by Saul, who describes a keen gardener uttering (6) in order to implicate (falsely) that she has had a great crop of tomatoes:

(6) I've got tomatoes coming out of my ears.¹⁷

Saul notes that we would intuitively judge this utterance to be a lie, although she then settles for a view of lying that is neutral with respect to whether or not metaphorical lies are possible. But it seems quite plausible that the gardener has indeed lied by uttering (6). For instance, she retains none of the deniability that misleaders typically retain with respect to the disbelieved content put forward: she could not (consistently) respond to accusations of lying by pointing out that she merely claimed that she had tomatoes coming out of her ears.¹⁸ Here are two further examples of non-literal lies that I have discussed in previous work:

Hyperbole

Carl desperately wants Daisy to come to his party, which is in full swing. When Carl calls Daisy, she says she'll come, but only if there is some food. Carl is aware that all the food has been eaten, but nonetheless utters:

[7] There's tons of food left.
Conversational implicature: There's lots of food left.

Irony

Greta and Henry are about to take a school exam. Greta has diligently prepared for the exam, but is aware that it would be decidedly uncool to admit this. When Henry asks whether she has studied for the exam, she rolls her eyes and utters:

[8] Of course I have studied for the exam.
Conversational implicature: I have *not* studied for the exam.¹⁹

¹⁷ Saul, *Lying, Misleading, & What is Said*, at p. 16.

¹⁸ It has been argued that such deniability-judgements can be used to decide whether an utterance is a lie or merely misleading. Cf. Andreas Stokke, "Lying and misleading in discourse," *Philosophical Review* 125 (2016): 83–134, at pp. 89–91, and Emanuel Viebahn, "Non-literal lies," *Erkenntnis* 82 (2017): 1367–1380, at p. 1370.

¹⁹ Viebahn, "Non-literal lies," at pp. 1368–1375. For further examples of and arguments in favour of non-literal lies, see David Simpson, "Lying, liars and language," *Philosophy and Phenomenological Research*, 52 (1992): 623–639, and Marta Dynel, "Comparing and combining covert and overt untruthfulness: on lying, deception, irony and metaphor," *Pragmatics & Cognition* 23, 1 (January 2016): 174–208.

In these cases, too, the speakers apparently do not retain deniability with respect to the conversational implicatures put forward, and intuitively it seems that they are lying, not merely misleading, about the content of the implicature.²⁰

There are thus several reasons to hold that it is possible to lie non-literally, and thus to hold that the lying-misleading distinction does not align with the distinction between asserting and conversationally implicating. Note that whether or not a lie is told literally or non-literally does not seem to make any difference for its moral evaluation: the non-literal lies (6), (7) and (8) appear to be wrong in exactly the same way as their literal counterparts are. This precludes a special treatment or quarantining for non-literal lies: these should receive the same treatment as literal lies.

From these two premises we get to the conclusion that at least some lies, namely those that are told with the help of conversational implicatures, do not involve the metasemantic unfairness that Berstler associates with lying. If it is possible to lie without misusing metasemantic conventions, then such conventions cannot ground a general moral advantage for misleading over lying.

Before moving on, I would like to consider two possible responses on Berstler's behalf. Firstly, one might think that Berstler could respond by pointing out that the semantic content of the examples of non-literal lies is false.²¹ For this reason, the objection continues, the speakers still misuse conventions of assertion as they assert the false semantic content of the sentence uttered. This could then be an interesting difference to standard examples of misleading, in which the semantic content of the sentence put forward is true.

However, such a response is unpromising because the speakers in the examples clearly do not assert the false semantic content of the sentences they utter. The gardener in Saul's example clearly does not put forward or even commit herself to the content that tomatoes are coming out of her ears; and the same holds, *mutatis mutandis*, for utterance (7). In uttering these sentences, the speakers are thus not misusing conventions of assertion.

Secondly, it might be argued that Berstler could accept the possibility of non-literal lies and could hold that such lies still involve the misuse of conventions in the following way: conventional language is used to assert something that is believed to be false (even if the assertion is non-literal). For

²⁰ For those worried that the non-literality in the first two examples has been conventionalised, here are some alternatives: instead of (6), the gardener might utter "The tomato god has been watching my crop," and still put forward the same implicature; and the implicature of (7) could equally be achieved by uttering "There's enough food to feed the entire town."

²¹ This is the case for examples (6) and (7). While the semantic content of (8) is true, this example is probably more controversial than the others, so I will focus on the former two examples for now.

example, in uttering (7), Carl uses the conventional meaning of “There’s tons of food left,” to assert something he believes to be false.

While I agree that liars do plausibly misuse conventions in this way, I think that this observation is of no help to Berstler’s argument in the current context, as the Lewisian metasemantic conventions, which are the ones that matter here, apply to *conventional* assertion only. As a result, Carl is not misusing a Lewisian convention if he utters (7) and thereby *non-literally* asserts something he believes to be false. Of course, it is quite plausible that there are conventions that govern non-conventional communication in general and non-conventional assertion in particular. Even with non-conventional communication, there is an expectation that speakers are sincere. But such conventions are not metasemantically significant on the Lewisian picture. And if they were, they could not set apart non-literal lies and cases of misleading, as in both kinds of cases conventional meaning is used to non-conventionally communicate something the speaker believes to be false.

It is thus far from clear that the lying-misleading distinction does align with the distinction between asserting and implicating, as Berstler assumes. Accordingly, it is also unclear whether fairness-related considerations always favour misleading in situations in which a lie is compared with a corresponding misleading utterance: with respect to metasemantic fairness, non-literal lies appear to be on a par with their merely misleading counterparts.²²

5. Lying, misleading and commitment

Finally, I would like to question the view that lying and misleading differ *only* in terms of fairness and are otherwise morally identical. To support this view, Berstler argues against existing proposals for morally relevant differences between lying and misleading and emphasises that assertion and implicature “play identical communicative roles” (WGL: 25). This leads her to conclude:

If there is a further moral difference between lying and misleading, it is incumbent on my opponent to state what it is. (WGL: 26)

In this section, I will take up Berstler’s challenge and point to a morally relevant difference between lying and misleading: a difference in communicative commitment. First, I will argue that lying and misleading differ in terms of communicative commitment. Then, I will try to show that the difference in commitment can be morally relevant.

²² Many thanks to an anonymous associate editor for comments that helped me to clarify this section.

To observe the difference in commitment between lying and misleading, let us return to the aforementioned example of the dying woman, this time in the original version by Saul.²³ In the example, a dying woman asks a doctor whether her son is well. The doctor saw the son yesterday, when he was fine, but knows that he was killed shortly afterwards. The doctor could then lie by uttering (9) or utter the true but misleading (10):

- (9) He's fine.
- (10) I saw him yesterday and he was fine.

According to Berstler, these two utterances play the same communicative role. But that is not quite right. While both utterances are used to put forward the same content, namely that the dying woman's son is fine, they do so in a different manner. In particular, (9) puts forward the relevant content in a *stronger way* than (10): with (9), the doctor takes on a communicative commitment to the son being fine, while she does not take on such a commitment with (10).

This view can be supported by considering how the doctor could react to potential challenges to her utterance. For example, having uttered (9), the doctor might be challenged to justify why she knows that the son is fine ("How do you know he is fine?"). It seems that the doctor cannot consistently dismiss such a challenge – she either has to provide the requested information or take back the content she put forward. By contrast, the doctor *can* consistently dismiss a challenge of that kind if she utters (10). In that case, she can point out that she did not claim that the dying woman's son is well, but merely that she saw him the previous day, when he was fine.²⁴

This difference in how speakers can react to challenges seems to support the view that lying and misleading differ in terms of commitment. Put differently: If we are comparing cases of lying and misleading, and if we hold everything fixed that does not have to do with the choice between lying and misleading, there can be a difference in commitment. But is the difference in commitment a *general* difference between lying and misleading? I think it is, and have argued for this position in

²³ Saul, *Lying, Misleading, & What is Said*, at p. 70.

²⁴ Berstler (WGL: 27–28) notes that misleading, but not lying, allows for certain denials of what one has done. In particular, misleaders have more ways of denying that they *meant* the content they insincerely put forward. I would like to note that such meaning-denials differ from the consistent dismissals I have just discussed, which concern the question whether the speaker has taken on a justificatory responsibility to a certain content, not whether a content was meant. I agree with Berstler that the possibility of meaning-denials is not relevant for morally evaluating choices between lying and misleading.

previous work.²⁵ While I cannot provide a full defence of that view here, I want to discuss an observation that might be taken as a reason to doubt that lying and misleading do in general differ in terms of commitment.²⁶ While it would be natural to challenge the doctor’s utterance of (9) with “How do you know he is fine?”, similar challenges would be strange or infelicitous in the examples Berstler discusses. For instance, following Judith’s utterance of (1) (“It’s delicious!”) in the initial example, it would be strange for Quentin to challenge Judith to justify how she knows that the cake is delicious – after all, he just saw her try it.²⁷ Do such cases undermine the commitment view?

There are reasons to believe that they do not. Firstly, the intuitive commitment-based difference between lying and misleading is present in Berstler’s examples, too. Intuitively, Judith commits herself to finding the cake delicious by uttering (1) but not by uttering (2) (“This is the best dessert I’ve had all week”). Secondly, there is a straightforward explanation for why challenges to justify knowledge are strange in Berstler’s examples: it is strange to ask for such a justification if the speaker’s purported evidence is obvious (even if the speech-act in question is committal). For the same reason, it would be strange to challenge a speaker who has asserted that she is in pain to justify how she knows she is in pain; here, too, we would not want to deny that the speaker has incurred a communicative commitment, even if we cannot rely on reactions to possible challenges to verify this. The observation about Berstler’s examples thus illustrates that there are some cases in which reactions to possible challenges cannot serve as a guide to commitment, but it does not show that lying and misleading need not differ in terms of commitment.

The question, then, is whether this difference in commitment can be morally relevant. Felix Timmermann and I have argued that it can be.²⁸ On our view, the lack of commitment can make misleading morally better in cases in which the speaker is uncertain whether or not the addressee wants an honest answer to a question. In such cases, the lack of commitment can leave a path to the truth. The addressee can detect the lack of commitment and can ask for a more committal answer. However, we also argue that the lack of commitment can make misleading worse in other cases, namely if there is an obligation for the speaker to leave the addressee in the dark about a certain matter. For instance, in the above example of John, the willingly deceived cook, it seems that Joanne

²⁵ See Emanuel Viebahn, “Lying with Presuppositions,” *Noûs* 54 (2020): 731–751, and “The Lying-Misleading Distinction: A Commitment-Based Approach,” *The Journal of Philosophy* (2021): 289–319.

²⁶ Many thanks to an anonymous associate editor for helpful comments on this point.

²⁷ The same holds for Berstler’s examples *Paper* and *Wife* (WGL: 10–11), where in each case the agent’s purported evidence is obvious.

²⁸ Timmermann and Viebahn, “To lie or to mislead?”, Section 4.

has an obligation to leave John in the dark about the quality of the soup, and this makes her misleading utterance (4) (“My mother used to cook that soup on Christmas Eve. I loved it.”) worse than her lie (3) (“It was the best fish soup I’ve ever tasted. It was delicious!”) – after all, the misleading utterance might raise John’s suspicion and thus reveal Joanne’s real opinion about the soup, which he does not want to hear.

We do accept that even if there is a difference in commitment between lying and misleading, this will not always be morally relevant. In particular, it will not be morally relevant in certain cases in which the addressee is not in a position to detect a lack of commitment.²⁹ But what matters in the present context is that there are *at least some cases* in which commitment is morally relevant. Berstler argues that lying and misleading are “morally identical except insofar as lying is unfair” (WGL: 7). To challenge this claim, it is enough to show that there are at least some cases in which lying and misleading are morally non-identical for reasons other than fairness. A difference in commitment is such a non-fairness-based reason, insofar as it can leave a path to the truth (and thus make misleading better in cases in which the speaker is uncertain whether or not the addressee wants an honest answer to a question) or can raise suspicions (and thus make lying better in cases in which there is an obligation to leave someone in the dark).

It is not our aim to show that commitment is the only factor that can set lying and misleading morally apart, and we explicitly leave open “the possibility that there are *further* features that can make for a moral difference between lying and misleading”.³⁰ So it is quite possible to combine our view with Berstler’s and to hold that commitment *and* fairness can bear on the moral evaluation of a choice between lying and misleading. But, of course, on such a combined view fairness will not be the only morally relevant factor, and there will be cases in which misleading is worse than lying (*contra* Berstler’s argument), namely cases in which the fairness-based advantage of misleading is outweighed by a commitment-based advantage of lying. The case of John, the willingly deceived cook, may well be an example of this kind.

If the foregoing is right, then lying and misleading not only play different communicative roles by differing in commitment, but this difference in commitment can also be morally relevant. While a fuller discussion of the apparent commitment-based difference between lying and misleading and also of its apparent moral significance is needed, initial impressions do suggest that fairness is not the only factor that can set lying and misleading morally apart. At the very least, Berstler’s claim that

²⁹ Ibid., 1495.

³⁰ Ibid., 1500. Original emphasis.

due to the unfairness of lying we “always have one extra reason not to lie” (WGL: 7) requires further support.

6. Conclusion

I have pointed to three difficulties for Berstler’s account. Firstly, it seems that intuitions do not straightforwardly support a view on which there is a general preference for misleading over lying; and in cases in which there is an intuitive preference for misleading, there also appear to be intuitions about *why* misleading is better that are at odds with Berstler’s account. Secondly, there are reasons to believe that the lying-misleading distinction does not align with the difference between asserting and implicating, in which case metasemantic fairness cannot ground a general preference for misleading over lying. Thirdly, Berstler does not do enough to establish that fairness is the only factor that can set lying and misleading morally apart. There appears to be a commitment-based difference between at least some cases of lying and corresponding misleading utterances, and this difference in commitment may well be morally relevant.

I am aware that none of these objections is unassailable. But I do think that, taken together, they are a significant challenge for Berstler’s account, partly because the objections are independent from each other: Berstler would have to show that all three objections are unsuccessful to fully restore her account.

At the same time, I want to highlight that I have not objected to Berstler’s central idea that literal assertion and conversational implicature, and thus many cases of lying and misleading, differ in terms of fairness. I think this is an important idea that should be explored further. Even if fairness does not give us a reason to favour misleading in every case, and even if fairness is not the only factor that can morally set apart lying and misleading, it does appear to be one factor that can bear on the moral evaluation of a choice between lying and misleading.