

# Assertion, Lying, and Untruthfully Implicating

Jessica Pepp

Uppsala University

*(This is a draft of a chapter for The Oxford Handbook of Assertion, ed. S. Goldberg, Oxford University Press)*

## 1. Introduction

Suppose that Gloria recently lost her job due to a personality clash with her boss. Suppose also that she is highly regarded in her field and already knows of several firms interested in hiring her, so she is not worried about her future employment situation. Gloria's father, however, worries a lot, and she suspects that knowing about her situation would cause him great unease. When he asks (as he often does) about her employment situation, she might reply, "I still have a job, don't worry," thereby lying to him. Or she might reply, "The company made a lot of cuts recently, but I wasn't affected by them." This would be a truthful assertion (assuming the company did make cuts, and assuming it was not the cuts that affected Gloria, but the personality clash). However, in saying this Gloria would *implicate* that she still has a job, which she believes not to be the case. She would untruthfully implicate, but she would not lie.

The difference between lying and untruthfully implicating is often taken to be morally significant. Contemporary philosophers often claim that people have a preference for untruthfully implicating over lying, and that people tend to believe that lying is morally worse than untruthfully implicating (other things being equal).<sup>1</sup> One study in behavioral economics provides some empirical evidence for this preference (though not specifically for the belief in moral difference).<sup>2</sup> Prominent philosophical and religious traditions treat avoidance of lying by resort to untruthful implicating as morally commendable.<sup>3</sup> And I would guess that most of us can recall occasions on which we have carefully avoided lying

---

1. See for example: Jennifer Saul (2012), Andreas Stokke's (2013) review of Saul, Melissa MacAulay and Robert Stainton's (2013) review of Saul, Jonathan Adler (1997), Roderick Chisholm and Thomas Feehan (1977), and Alan Strudler (2010).

2. See Uyanga Turmunkh, Martijn van den Assem, and Denni van Dolder (2017).

3. For instance, the Kantian and Augustinian traditions. See Macintyre (1995) for discussion.

by implying something we do not believe rather than coming out and asserting it.

But while morally relevant similarities between lying and untruthfully implicating are readily apparent (for instance: both are typically deceptive, manipulative, or harmful to one's interlocutor), morally relevant differences are harder to identify. The intuition that the difference between lying and untruthfully implicating is morally significant—what I will call the "Difference Intuition"—calls out for justification.

This chapter explores the prospects for justifying the Difference Intuition. I begin, in section 2, by getting clearer on the difference between lying and untruthfully implicating, and the role that assertion plays in making this difference. Then, in section 3, I set out various ways of sharpening the Difference Intuition. In section 4, I survey a range of approaches to justifying one class of sharpenings of the Difference Intuition. In section 5, I sketch an approach to justifying an alternative sharpening of the Difference Intuition, which is inspired by John Stuart Mill's discussion of lying. Section 6 concludes.

## 2. Assertion, Lying and Untruthfully Implicating

Lying is often defined in terms of assertion, along the following lines:

To lie is to assert that something is the case while believing that it is not the case.

This definition of lying requires a notion of assertion that rules out asserting something ironically, in jest, in a play, metaphorically, by accident, or in any other way that would allow for asserting something one believes is not the case without lying. Different notions of assertion might be imported into the definition of lying in order to do this work.<sup>4</sup>

Alternatively, lying can be defined in terms of a thinner notion of assertion, with an additional requirement to prevent instances of speaking ironically, jokingly, and so on from counting as lies. For instance, the third clause of Jennifer Saul's definition of lying does this work:

If the speaker is not the victim of linguistic error/malapropism or using metaphor, hyperbole, or irony, then they lie iff (1) they say that P; (2) they believe P to be false; (3) they take themselves to be in a warranting context.<sup>5</sup>

---

4. For instance, Stokke's (2013) definition of lying treats asserting as proposing to add information to the conversational common ground. Keiser (2016) treats assertion as the expression of belief, where expression of a belief involves providing one's audience with a reason to believe one has that belief.

5. Saul (2012: 3).

To take oneself to be in a warranting context is to take oneself to be in a context in which one implicitly guarantees the truth of what one says (unlike in contexts of joking or acting).<sup>6</sup> Hence, one way to view Saul's definition of lying is as a definition based on a more substantive notion of assertion, according to which assertion involves making a certain kind of commitment or guarantee.

The literature on how lying should be defined is expansive, offering many candidate definitions to choose from.<sup>7</sup> I will just mention two features of the basic definition I gave above that could be varied. First, the definition does not require that what one asserts is not the case (but only that one believe that it is not the case). Second, the definition does require that one believe that it is *not* the case (instead of only requiring that one not believe that it is the case). It would not substantially affect the discussion of this chapter to make the opposite choice on either of these features.

By requiring that a speaker assert that something is the case, most definitions of lying differentiate it from various other ways that one might convey that something is the case while believing that it is not the case. The verb "implicate" and the noun "implicature" are used to refer to various ways of *meaning*<sup>8</sup> that something is the case without asserting that it is the case. Implicature is often divided, following Paul Grice,<sup>9</sup> into conventional and conversational implicature. Conventional implicature, as in an assertion of, "He is an Englishman; he is, therefore, brave," is determined by the linguistic conventions associated with the words used; conversational implicature, as in the example of Gloria in the previous section, is determined by conversational norms. Mitchell Green points out that there is also room for a notion of non-conventional, non-conversational implicature, as when one means that one dislikes a film by making a sour face at the mention of its title.<sup>10</sup> With any type of implicature, one may implicate that something is the case while believing that it is not the case: one may untruthfully implicate.

Just as, on a thin notion of assertion, one might assert in a play, in jest, ironically, and so on, one might also implicate in these ways. To play on one of Grice's famous examples

---

6. Saul follows Carson (2006) in using this notion of warrant to define lying.

7. For a thorough taxonomy of definitions of lying and considerations for and against the various definitions, see Mahon (2016).

8. Here I use "meaning" in the sense that Grice (1957) called "non-natural meaning." It is now often called "speaker meaning." I am not going to offer or endorse any particular account of speaker meaning. I only point to it as something speakers do. There are ways of conveying that something is the case without meaning that it is. For instance, one might act as if something is the case.

9. Grice (1984).

10. Green (2007: 100).

of implicature, suppose that Smith objects fervently to the existence of the university cheerleading squad and has been spending his days protesting outside their clubhouse. If A says to B, "Smith doesn't seem to have a girlfriend these days," and B replies "He has been paying a lot of visits to the cheerleading clubhouse lately," then B implies that Smith has a girlfriend on the cheerleading squad, but only in jest. A and B both assume that Smith would never date a cheerleader. But what B does here is not the kind of untruthful implicating that is to be compared with lying. Rather, the acts of implicating that constitute the comparison class of untruthful implicature must fulfill the same sort of conditions (expressing a belief, proposing to add information to the common ground, committing to the truth of what is implicated, or the like) as the acts of asserting that constitute lying.<sup>11</sup>

One further complication in comparing lying and untruthfully implicating is due to the fact that implicating often consists in giving reasons for drawing some conclusion, without stating the conclusion. In such cases, often what is implicated is that something may be, or is likely to be, the case, rather than that it *is* the case. Consider Grice's original girlfriend case:

A: Smith doesn't seem to have a girlfriend these days.

B: He has been paying a lot of visits to New York lately.

(Grice 1989: 32)

As Grice put it, "B implicates that Smith has, *or may have*, a girlfriend in New York." (emphasis added) How hedged B's implicature is might depend on aspects of his delivery, such as whether he smiles knowingly or shrugs uncertainly. That some cases of implicature are hedged might make it seem as though implicating in general signals uncertainty, as compared with asserting. This in turn might suggest that untruthfully implicating is morally preferable to lying because it gives the hearer fair warning not to rely too much on what is implicated. But implicating does not in general signal uncertainty: consider Grice's example of the reference letter writer who writes, "Mr. X's command of English is excellent and he always comes on time to tutorials," or Jennifer Saul's example of Bill Clinton's famous assertion, "There is no sexual relationship." These speakers do not implicate that Mr. X *might* be bad at philosophy or that there *might* never have been a sexual relationship, they implicate these things in an unhedged way. If these implicatures are morally preferable to the corresponding outright lies, it is not because the former indicate some sort of uncertainty or hedging while the latter do not. In cases of untruthfully implicating where

---

11. We will also have to decide whether what is implicated must not be the case, or must only be believed not to be the case, and whether it must be believed not to be the case or only not believed to be the case. Again, I do not think these choices will matter for the present discussion.

the implicating is hedged (for instance, if B in Grice's case believed that Smith does not have a girlfriend in New York), the relevant comparison would be with a case in which the corresponding assertion is similarly hedged (for instance, if B asserted that it might be that Smith has a girlfriend in New York<sup>12</sup>).

To summarize: the distinction between lying and untruthfully implicating whose moral significance is at issue is the difference between asserting (where assertion is understood in some sense that rules out joking, speaking ironically, etc.), in an unhedged way, that something is the case although you believe that it is not the case, and meaning (also understood in a sense that rules out joking, etc.), also in an unhedged way, without asserting, that something is the case although you believe that it is not the case. (As suggested above, readers should feel free to add the requirement that what is asserted in fact *not* be the case, or to substitute the requirement that what is asserted *not be believed to be the case* in place of the requirement that it be *believed not to be the case*.)

### 3. The Intuition of Moral Difference

So far, I have formulated the Difference Intuition loosely, as the intuition that the difference between lying and untruthfully implicating is morally significant. This intuition might be made more precise, so as to be a suitable target for justification, in different ways.

One way is to make a universal claim about pairs of contrast cases. One might claim that whenever two (possible) cases differ only in that one is a case of lying that *p*, while the other is a case of untruthfully implicating that *p*, the act in the first case is worse than the act in the second case. This seems to be the approach of Saul (and others in the recent debate surrounding her book), who sharpens the Difference Intuition to:

(M) Holding all else fixed, lying is morally worse than merely deliberately attempting to mislead; and successful lying is morally worse than merely deliberately misleading.<sup>13</sup>

---

12. I am not sure whether we would intuitively count such a hedged untruthful assertion as a lie or not.

13. (Saul 2012: 72). Although Saul contrasts lying with attempting to mislead, rather than with untruthfully implicating, her focus seems to be almost entirely on attempting to mislead by untruthfully implicating. Indeed, she suggests that attempting to mislead by acting as though something is the case, rather than by *meaning* that it is the case, might well be morally better than lying and at any rate is not the kind of "linguistic deception" with which she is interested in comparing lying. (75-77)

A further note about Saul's (M) is that attempting to mislead (or deceive) by untruthfully implicating is not obviously the appropriate contrast with lying. There is a growing consensus that not all lies are attempts to mislead or deceive the audience. (See the discussion of bald-faced lies and other forms of non-deceptive lying in Carson (2006), Fallis (2009), and Sorensen (2007).) Accordingly, the moral significance of lying may not depend on its being a means of deceiving. (See Stokke, forthcoming and Shiffrin 2014 for ethical accounts of lying that do not place deception at the center.) Given this, the relevant contrast to evaluate for moral

It is clear from the counterexamples to M that Saul adduces that she intends M as a universal claim across pairs of contrast cases. One such counterexample is a pair of cases in which a murderous chef has cooked dinner in peanut oil for his guest, who is deathly allergic to peanuts. In one case, when asked whether the dinner contains peanuts, the murderous chef replies "No, I didn't put any peanuts in," and in the other case he replies, "No, it's perfectly safe for you to eat." Saul's intuition is that the chef's act of lying in the second case is no worse than his act of untruthfully implicating in the first case. She takes this to show that M is false.

Saul recognizes that such counterexamples could be parried by arguing that although there is a moral difference between the cases, the moral status of both acts is so extreme that our intuition is unable to pick up the difference. But she also proposes as a more plausible sharpening of the Difference Intuition the following thesis:

(M-D) Except in certain special cases: holding all else fixed, lying is morally worse than merely deliberately attempting to mislead; and successful lying is morally worse than merely deliberately misleading.<sup>14</sup>

This thesis allows for some contrast pairs in which the lying/attempting to mislead (i.e. untruthfully implicating) difference does not make for a difference in the moral status of the two acts.

One might also think that although the difference between lying and untruthfully implicating is morally significant, it need not typically or most of the time make a difference in the moral status of contrast pairs, but need only do so in some cases. (As Shelly Kagan has argued, it is not obvious that a difference between two ways of behaving can make a difference in the moral status of *some* pairs of contrast cases only if it makes a difference in the moral status of *all*—or even most—contrast pairs.<sup>15</sup>) Hence, one might sharpen the Difference Intuition as follows:

At least for some pairs of cases: holding all else fixed, lying is morally worse than untruthfully implicating.

Finally, the Difference Intuition might be sharpened without appeal to pairs of contrast cases. One might claim that when faced with a need to convey that something is the case despite believing that it is not, it is good moral practice to untruthfully implicate rather than lie. This might be so even if there are no pairs of contrast cases in which the

---

significance is not the contrast between lying, on the one hand, and attempting to mislead or deceive (without lying), on the other, but between lying and untruthfully implicating, regardless of deceptive intent.

14. Saul (2012: 74).

15. Kagan (1988: 12-14).

case of lying is morally worse than the case of untruthfully implicating. (For instance, a maximizing act consequentialist might hold that no lying/untruthfully implicating contrast pairs display a difference in moral status, given that total outcomes of the relevant acts are held fixed across the cases. Nonetheless, they might advocate a rule of thumb or decision procedure favoring untruthfully implicating over lying.<sup>16</sup>)

I mentioned in the Introduction that many philosophers claim to have the Difference Intuition (in one form or another), and that it seems to them to be widespread. However, I am not aware of any clear evidence that the intuition is (or is not) widespread. Anecdotally, I find that most people I ask do acknowledge at least a weak form of the intuition, but some say that they cannot see any moral difference, or even that they have the opposite intuition, that untruthful implicating seems morally worse than lying.<sup>17</sup> This "Reverse Difference Intuition" could be sharpened in all the same ways as I have described for the Difference Intuition.

#### **4. Approaches to Justifying the Difference Intuition**

The first three sharpenings of the Difference Intuition that I discussed in the previous section claim, respectively, that an act of lying always has, typically has, or sometimes has a different moral status from the corresponding act of untruthfully implicating, assuming everything else is held fixed. In this section, I will discuss three broad approaches to justifying such claims. The first approach is to try to show that there are different morally relevant factors in cases of lying as compared with cases of untruthfully implicating. The second is to try to show that although the morally relevant factors are of the same kind in both cases, they differ in degree. The third approach eschews the weighing of morally relevant factors in determining an act's moral status, aiming instead to identify fundamental moral principles that lying, but not untruthfully implicating, violates. I will explore the prospects for each of these broad approaches in turn.

##### **4.1. Different morally relevant factors**

One tempting suggestion is that doing harm is a morally relevant factor in (at least many cases of) lying, but not in untruthfully implicating, where one merely allows harm.<sup>18</sup>

---

16. I discuss justification of the Difference Intuition along these lines in section 5.

17. These intuitions support the argument of Clea Rees (2015), that in ordinary conversational settings lying is morally better than attempting to mislead by untruthfully implicating.

18. Both Jonathan Adler (1997: 446) and Bernard Williams (2002: 101) observe a surface similarity between

It might be thought that because the untruthful implicater leaves the hearer to draw her own inferences, the communication to the hearer that something is the case which the untruthful implicater believes is not the case is not really something the untruthful implicater *does*. But this is not right: communicating that something is the case is exactly what an implicater attempts to do in implicating that it is the case, just as much as when she asserts outright that it is the case. Either way, some participation on the hearer's part is required, as hearers must make inferences in order to understand assertions (e.g. they must make inferences about what names, indexicals, and ambiguous words refer to). And in both kinds of case, these inferences are often quite automatic.<sup>19</sup> To whatever extent one *does* harm by lying on a given occasion, one would also *do* harm by the corresponding act of untruthfully implicating (holding other things fixed).

A different way to make a case for a difference in the morally relevant factors in lying and in untruthfully implicating is to appeal to the nature of assertion. On some views, part of what makes an act be an assertion that something is the case is the fact that the speaker incurs certain responsibilities for its being the case.<sup>20</sup> These responsibilities are not usually taken to be moral responsibilities, but rather normative responsibilities of some other stripe, such as epistemic responsibilities or perhaps just conversational responsibilities.<sup>21</sup>

In this vein, Jonathan Adler argues that the practice of assertion is governed by a "demand of truthfulness." He illustrates this demand by observing that it is fair to challenge an assertion by asking the asserter how she knows or what her evidence is. By contrast, challenges like this are not appropriate for implicatures, indicating that our linguistic practice does not place the same "demand of truthfulness" on implicating. Although the norms of our linguistic practice are not themselves moral norms, they have "moral force" or "nor-

---

the lying/untruthfully implicating distinction and the doing/allowing distinction.

19. Saul makes both of these points. (2012: 80-81)

20. Examples include Searle (1969), Brandom (1994), and Rescorla (2009). Saul's and Carson's appeal to warranting in their definitions of lying also suggest a view of assertion along these lines. (See note 6.)

21. One exception is Terence Cuneo's (2014) view of speech acts. According to Cuneo, what makes an action be an assertion that something is the case is, in part, the fact that the agent acquires *moral* responsibilities for things being as she asserts that they are. Cuneo does not discuss implicature, but one can imagine a justification of the Difference Intuition within his framework along the following lines: Part of what it is to assert that something is the case is to acquire a moral responsibility for it being the case. Acquiring a moral responsibility for something's being the case is not part of what it is to implicate that it is the case. Hence, lying violates a moral responsibility that untruthfully implicating does not violate. However, I think such a rationale would be problematic for parallel reasons to those I am about to adduce against Adler's justification. Namely, it is not clear why an advocate of Cuneo's view should deny that what it is to implicate, at least in the "thick" sense with which we are concerned (see Section 2) and which parallels Cuneo's notion of assertion, is in part to acquire moral responsibilities for what one implicates being the case. Of course, one could insist that this responsibility is lesser or absent in the case of implicature, but this would simply be to restate the Difference Intuition in the context of the theory, not to justify it.



mative weight" because they are "mutually expected, rational, and tacitly endorsed."<sup>22</sup> Failing to abide by such norms is morally significant—perhaps because it is a violation of fairness.<sup>23</sup> Since there is such a norm that demands (inter alia) that one believe what one asserts, but no such norm that demands that one believe what one implicates, lying is morally significant in at least one way that untruthfully implicating is not.<sup>24</sup>

A similar suggestion to Adler's is made by Matthew Benton, who derives a norm against lying from the "epistemic norm of assertion", by which he means a "social linguistic rule" that specifies a necessary condition for assertion to be epistemically proper.<sup>25</sup> Given that some such rule is in place, asserting that something is the case while believing that one does not meet the specified requirement is epistemically "vicious."<sup>26</sup> However, according to Benton, "one does not undertake the same epistemic responsibility by conversationally implicating that *p* as when one asserts that *p* with full force."

Whatever the merits of the idea that failing to abide by mutually endorsed rules of conversation is morally significant, it is not clear that such rules in fact treat asserting and implicating so differently. To demonstrate that they do, Adler appeals to the fact that implicature can be used to make available information that one is not prepared to fully defend, as when one offers one's evidence for a conclusion, intending the hearer to draw the conclusion herself. Adler compares this to saying, "It may be that *p*." But, as discussed in section 2, it is not the fact that one is implicating rather than asserting that makes such implicatures have a hedged character. Rather, other features of the context and manner of expression accomplish this. Implicatures are not in general hedged, so it is not clear why the demands of truthfulness for implicature would in general be less.

Adler and Benton also appeal, respectively, to the inappropriateness of challenges to implicatures and to the availability of replies like "I never said that." But while it might be inappropriate to challenge an implicature by asking "Why do you think that?" or "How do

---

22. Adler (1997: 449-51).

23. Adler does not defend the move from the norms of some practice being mutually known and endorsed to their having moral significance, but he does refer to the differential norms of asserting and implicating as a "fair understanding." For a defense of the move from mutually manifest norms (of the practice of assertion, specifically) to moral significance, see Goldberg (2015: 175-80).

24. Stuart Green seems to take a similar line, though he puts the point in terms of a "warranty of truth" being given by an asserter but not by an implicater. (Green 2006: 79)

25. Candidate conditions include knowing that what one asserts is true, believing it, or having sufficient evidence for it.

26. Benton (forthcoming). It is not entirely clear whether Benton takes this epistemic viciousness to be *morally* significant. He says that the viciousness is "primarily epistemic," and thus exhibited even by lies "undertaken with high moral or social aims." But he also says that when people rely on lying testimony, "it's not just that one's epistemic position is compromised; the social goodwill we have toward others, especially those with whom we have ongoing relationships, is strained."

you know that?", similar challenges that acknowledge how the implicature was made are appropriate. Here are some examples:

i. Girlfriend-in-New-York example:

A: Smith doesn't seem to have a girlfriend these days.

B: He's been paying a lot of visits to New York lately.

(Imagine this said archly, so as not to suggest hedging or uncertainty.)

A: Why do you think that means he has a girlfriend there?

ii. Deathbed question example:

Woman on her deathbed: Is my son alright?

Kindly companion: I saw him yesterday and he was just fine.

Woman on her deathbed: How do you know nothing terrible happened to him today?

iii. Marriage and children example:

A: John and Jane got married and had a baby.

B: How do you know it happened in that order?

These interactions suggest that the demand of truthfulness (or for appropriate epistemic authority) applies to implicatures as well. It is true that the challenged speaker in each example can reply, "I didn't say that." But this does not seem to respond to the challenge, which demands the speaker's grounds for implicating, not her grounds for asserting. The challenged speaker could respond more directly to the challenge by saying, "I didn't mean to imply that." But a speaker challenged on an assertion might equally respond with "I didn't mean to say that." Both responses seem to acknowledge a responsibility to defend the assertion or implicature that the hearer took the speaker to have made, if the speaker had in fact made it.

As far as I can tell, conversational norms no more allow for untruthfully implicating than they allow for lying. Thus, I do not see a viable defense of the idea that lying, but not untruthfully implicating, is morally significant in virtue of being a failure to play by the rules of a mutually endorsed conversational practice.

## 4.2. Same morally relevant factors, different degrees

I noted at the beginning of the chapter that morally relevant similarities between lying and untruthfully implicating are easy to identify. Accordingly, some philosophers think that the same factors are morally relevant across the two types of act, but that they differ in degree.

For example, Alan Strudler argues that what is inherently morally significant about both lies and untruthful implicatures is that they are attempts to breach the trust of another person. Whether one asserts that something is the case or implicates that it is the case, one invites the hearer to trust that one is a cooperative communicator, and thus that one believes that what one asserts or implicates is in fact the case. If one believes that what one asserts or implicates is not the case, one intends to violate the very trust one is inviting.<sup>27</sup>

Nonetheless, Strudler argues that in asserting that something is the case, one asks the hearer to trust one *more* than one would in implicating that it is the case. In asserting, but not in implicating, one asks the hearer to surrender the option to raise a doubt about what is asserted without sacrificing the trusting nature of the conversation. Here is an illustration of the contrast, using an adaptation of Strudler's central example:

a. Raising a doubt about what is implicated; no loss of trust

Prospective house buyer: This would be a great neighborhood for my kids.

House seller: So you intend to move into this house with your family?

b. Raising a doubt about what is asserted; loss of trust

Prospective house buyer: I intend to move into this house with my family.

House seller: Really, you do?

Strudler's claim is that the buyer asks the seller to trust him more—to surrender more of her options—by asserting, as in case b, than by implicating, as in case a. Thus, if the buyer does not intend to move into the house with his family, asserting this is a greater violation of trust than implicating it.

The problem with this rationale is that while raising doubts about what is asserted sometimes breaks down trust, and raising doubts about what is implicated sometimes does not, the reverse seems just as likely. Even with cases a and b, one can imagine it go-

---

27. Strudler (2010: 173-4).

ing either way. If, in case a, it is perfectly clear that the buyer is presenting himself as considering the house as a future home for his family, the seller's question will sound unduly suspicious and will likely break down trust. In case b, the seller's question might come off as suspicious, but it also might come off as surprised and curious—a request for some further explanation of what was asserted.<sup>28</sup> In general, it does not seem that assertions invite a greater surrender of trust-preserving conversational options than implicating. Hence, there is no reason to think that lying is a greater violation of trust.

Reversing Strudler's claim, Clea Rees argues that untruthfully implicating is generally a greater violation of trust than lying, at least in "contexts characterized by standard assumptions of communicative cooperation."<sup>29</sup> Granted, there are special contexts (courts of law, business negotiations, certain governmental forums, adversarial interviews, etc.) where it is expected that people will both aim to deceive if it benefits them, and aim to avoid lying. (Bernard Williams calls these contexts "adversarial but rule-governed,"<sup>30</sup> a label which I will adopt.) In such contexts, Rees allows, lying is generally morally worse than untruthfully implicating. But in standard contexts, she claims, trusting both the assertions and implicatures of our interlocutors is an epistemic and moral obligation. Thus, when we assert or implicate that something is the case, we generate an obligation on the part of our interlocutor to trust us. This means that in trusting our implicatures, our interlocutors need to exercise "the moral and epistemic goodwill required to presume [us] fully cooperative conversational partners," whereas to trust our assertions, they need muster only "the minimal trust involved in assuming others' assertions to be honest." Thus, in implicating, we obligate others to trust us more extensively, making it a more severe violation of this trust if we do not believe what we implicate.

Rees's reversal of the usual starting point that lying is intuitively morally worse is interesting. But the argument is problematic. Suppose Rees is right that in standard contexts hearers have an epistemic and moral obligation to treat speakers as fully cooperative conversational partners, trusting both their assertions and their implicatures. Then in such contexts the trust that is violated either by a lie or by an untruthful implicature is the same robust trust. It is true that this trust is greater than the trust that is violated by lying in ad-

---

28. Saul (2012: 78-80) criticizes Strudler's argument in a similar vein, although she focuses on whether lying and untruthfully implicating differ in their likelihood of breaking down productive exchange once the lie or untruthful implicature has been revealed. Strudler's point, though, is not that revealed lies are more destructive of productive exchange than revealed untruthful implicatures, but that lies (revealed or not) are greater violations of trust than untruthful implicatures.

29. Rees (2013: 64).

30. Williams (2002: 109).

versarial but rule-governed contexts. But this does not imply a difference *within* a standard context between the trust violated by untruthfully implicating and the trust violated by lying.

I will round out this section with a suggestion for justifying the Difference Intuition from the perspective of care ethics. Broadly speaking, care ethics takes the caringness of the attitude or motivation with which an act is performed as the central factor relevant to the moral status of the act. In one development of this kind of view, on which caring is understood in terms of empathy, Michael Slote argues that those who kill innocents "in cold blood," face to face, exhibit a greater lack of empathy than those who kill innocents at a distance (for example, from an airplane). On Slote's view, this makes killing in cold blood morally worse than killing at a distance.<sup>31</sup> Slote does not discuss the lying/untruthfully implicating distinction, but one can imagine a parallel rationale for seeing lying as morally worse. Putting it a bit figuratively: when you lie to someone, you hear (or see) quite immediately your presentation of things as being a way that you do not take them to be. You do not have to imagine your audience's processing of an implicature, and you cannot distract yourself by focusing on truthful assertions that you might use to effect an implicature. You are face to face with your delivery to your interlocutor of (what you take to be) a falsehood. If (and when) lying involves a lack of empathy for your interlocutor,<sup>32</sup> it might be that the corresponding untruthful implicature would involve a lesser lack of empathy, since it would not require as direct an awareness of the offense done to the interlocutor. Thus, it might be that on Slote's view untruthfully implicating is morally better.

The evaluation of this suggestion (which, again, is not one that Slote himself makes) seems to me to depend on answering partially empirical questions about empathy (for instance, is empathy involved in our apparent preference for lying over untruthfully implicating?). It also, and more critically, depends on assessing the view that empathy is a, or perhaps the, central morally relevant factor in determining an act's moral status. Both of these tasks are beyond the scope of this paper. I will rest with identifying this as a potential route for justifying the Difference Intuition.

---

31. Slote (2007: 25).

32. Slote has an account of the ways in which lying can display a lack of empathy, but the details are not important for the present point. (2007: 52-3)

### 4.3. Fundamental Moral Principles

Some normative ethical views are best seen as articulating fundamental moral principles, rather than as identifying and weighing up morally relevant factors. Kantian ethics is perhaps the central exemplar, especially in discussions of lying. Kant himself seems to have held that lying is worse than untruthfully implicating—indeed, that lying is absolutely forbidden while untruthfully implicating is often permissible.<sup>33</sup> It is less clear on what grounds he held this.

Alasdair Macintyre reads Kant as holding that "my duty is to assert only what is true and that the mistaken inferences that others may draw from what I say or what I do are, in some cases at least, not my responsibility."<sup>34</sup> Macintyre goes on to say that one might be tempted to think it is the non-universalizability of the maxim that tells one to lie whenever it is convenient that, for Kant, grounds the duty not to assert what one takes to be false. Non-universalizability might also be tempting as a tool for justifying the moral difference between lying and untruthfully implicating.<sup>35</sup> I am not aware of any previous attempts to do this, but how might one go? One way might be to argue that universalization of the maxim, "lie whenever it is convenient," would yield the maxim that everyone lies whenever it is convenient. But if everyone lied whenever it was convenient, no one would ever take seriously anything that was said, so it would be impossible to assert anything, and so impossible to lie. Hence, in acting on the maxim that she lie whenever it is convenient, the liar cannot simultaneously and without contradiction will that this maxim become a universal law. This "contradiction in conception" generates, in the Kantian framework, a perfect duty not to lie.

Suppose we grant this rationale for a perfect duty not to lie.<sup>36</sup> What would be the grounds for thinking that a duty against untruthfully implicating could not be generated in the same way? The untruthful implicater's maxim would be something like, "untruthfully implicate whenever it is convenient." The universalization of this would be that everyone untruthfully implicates whenever it is convenient. If it is right that universal lying-when-convenient would result in no one ever taking seriously anything that is asserted, it should also be plausible that universal untruthfully-implicating-when-convenient would result in no one ever taking seriously anything that is implicated. (This is some-

---

33. See Macintyre (1995: 336-7) for discussion.

34. Macintyre (1995: 337).

35. At least, whenever I discuss the Difference Intuition with philosophers, someone suggests that Kantians might use a difference in universalizability to justify it.

36. Pace Macintyre, who follows Korsgaard (1986) in rejecting universalizability considerations as grounding the duty against lying.

thing like what happens in the adversarial but rule-governed contexts discussed in the previous section.) Further, if it is right that the former result would mean that assertion, and thus lying, would become impossible, it seems equally plausible that the latter result would mean that implicature, and thus untruthful implicature, would become impossible.<sup>37</sup>

A different way of justifying the Difference Intuition on universalizability grounds might start by arguing that universalization of the liar's maxim yields a "contradiction in willing," rather than a contradiction in conception. A maxim is said to lead to a contradiction in willing if the agent cannot rationally will that it be universalized (even if the universalization does not defeat the agent's purpose in acting, as in a contradiction in conception). Seana Shiffrin suggests such a view of the wrong in lying:

...If universalized, [the liar's] maxim transforms a mechanism for exclusively conveying the truth into a mechanism for conveying both the false and the true. In this case, versatility is a vice. By doing so, one has eliminated a fail-proof, trustworthy mode of access to one another. Because there are no external methods of verification—no means by which others may peer into one's mind—to lie is to sully the one road of authoritative access to oneself and thereby cut oneself off from community with others. Doing so frustrates achievement of the compulsory moral ends associated with mutual understanding and cooperation.<sup>38</sup>

Shiffrin distinguishes the wrong of lying from the wrong of deceiving (which she takes to involve a contradiction in conception). She does not distinguish the wrong of lying from the wrong of untruthfully implicating, however.<sup>39</sup> On her view, it is speech, or linguistic communication, that constitutes the "one road of authoritative access" to each other's thoughts. But implicature is also part of linguistic communication, so untruthfully implicating would also seem to be acting on a maxim whose universalization would sully this critical tool. There does not seem to be anything in Shiffrin's account that would ground a moral difference between lying and untruthfully implicating.

One possibility would be to supplement Shiffrin's account with a suggestion from Jonathan Webber. Like Shiffrin, Webber emphasizes our collective need to be able to rely on each other's testimony and claims that lying damages this ability more than untruthfully implicating. The reason for this, according to Webber, is that having lied makes both one's assertions and one's implicatures less credible, whereas having untruthfully impli-

---

37. In this scenario it would still be possible to make and derive implicatures in one sense, but this would not be the robust kind of implicature that parallels the robust kind of assertion needed for lying. (See Section 2 above.) In the same way, in the scenario where the liar's maxim is universalized it might still be possible to say things in one sense, but not to engage in the robust kind of assertion needed for lying.

38. Shiffrin (2014: 41).

39. Recall that neither lying nor untruthfully implicating need be attempts to deceive. (See note 13.)

cated only makes one's implicatures less credible.

Incorporating this claim into Shiffrin's account of the moral significance of lying, it might be argued that the liar's maxim, if applied universally, would make it irrational to trust either what speakers assert or what they implicate. By contrast, the rule by which the non-lying untruthful implicater acts, if applied universally, would only make it irrational to trust what speakers implicate. This would still leave us with the ability to rationally trust each other's assertions, arguably ensuring the needed basis for moral life, even in difficult circumstances where "the behavioral components of trust have frayed."<sup>40 41</sup>

But why should we think that the liar's maxim, if applied universally, would make it irrational to trust either what speakers assert or what they implicate? Why would it not *only* make it irrational to trust what speakers assert, leaving what they implicate trustworthy enough?

One potential rationale, suggested by Webber's discussion, is as follows. If it is not rational to trust what speakers assert, then it is not rational to treat speakers as cooperative. If it is not rational to treat speakers as cooperative, then it is not rational to draw any implicatures from their conversational contributions, because implicature is based on the assumption that speakers are cooperating in furthering conversational ends. So once the rationality of trusting assertions is lost, there is no possibility of a respite in trustworthy implicatures.

This rationale is questionable. Suppose that the liar's maxim were universalized, so that it would not be rational to trust speakers' assertions (in the absence of specific grounds for trust). It is not clear why this should entail a general breakdown in cooperative communication. For instance, it might still be that people pose questions, make com-

---

40. Shiffrin emphasizes the importance of this ability for rational trust in difficult circumstances in a follow-on piece. (2016: 244).

41. Webber himself offers a different justification of the Difference Intuition that I confess to not grasping fully. He argues that because society needs reliable informants, and because liars compromise their own reliability more than untruthful implicaters, this "justifies society in reserving a more severe opprobrium for lying than is to be employed in response to misleading [defined here, roughly, as untruthful implicating]." Further:

Because it is right in this way for society to treat liars more severely than misleaders, each of us has good reason to mislead rather than to lie when faced with that choice. For each member of society has good reason to have a stronger negative attitude towards lying than towards misleading. Each of us needs to be able to rely on others as informants. We can still rely on misleaders, though we need to be careful, but cannot rely on liars. Each of us, therefore, has good reason to prefer to mislead than to lie, because according to the standards we ourselves ought to have, lying is more condemnable than misleading. (2013: 654)

One way of interpreting this might be along the lines of the Millian proposal discussed in Section 5 below.



mands, and express feelings sincerely. Further, it might still be that people implicate sincerely, even when doing so via the sullied speech act of assertion (at least, in cases where what is implicated is independent of what is asserted.<sup>42</sup>) It might remain perfectly rational to trust implicatures, despite its not being rational to trust assertions.

This seems far-fetched because of the central role assertion actually plays in our communicative practice. We have the sense that if you could not trust what others straight-out asserted, you could not trust them at all. Whereas, because we are familiar with Williams's "adversarial but rule-governed" contexts, we have the sense that there is still room for trusting assertions even when trust in implicatures is irrational. But there does not seem to be any reason in principle that a breakdown of credibility in assertion would entail a breakdown of credibility in implicature, while a breakdown of credibility in implicature would not entail a breakdown of credibility in assertion. Indeed, in ordinary contexts, when someone is seen to untruthfully implicate it is likely that their credibility in assertion *is* reduced, as well as their credibility in implicature.

In sum, it is not clear that there is a difference in universalizability between the liar's maxim and the non-lying untruthful implicator's maxim, either with respect to a contradiction in conception or a contradiction in willing.

## 5. A Millian justification of the Difference Intuition

All of the approaches to justifying the Difference Intuition just surveyed aim to show that the difference between lying and untruthfully implicating could make a difference to the moral status of an act, when all else is held equal. But I mentioned at the beginning of the chapter that one might reject this, but still uphold the Difference Intuition as shaping a proper moral decision procedure, or rule of thumb. In this section, I will explore such a justification, inspired by John Stuart Mill's discussion of lying.

For simplicity, let us interpret Mill as a maximizing act-utilitarian. So interpreted, Mill advocated a rule of thumb against lying on the grounds that our ability to trust each other's assertions is "the principle support of all present social well-being" and that the "insufficiency" of this ability to trust each other's assertions "does more than any one thing

---

42. To illustrate: if I lie in asserting "John and Jane got married and had a baby," then I also untruthfully implicate that John and Jane got married first and had a baby second, since in order to do those two things in that order you have to do those two things. On the other hand, if I lie in asserting, "He has an excellent command of English and always comes on time to tutorials," I may nonetheless truthfully implicate that he is no good at philosophy.

that can be named to keep back civilisation, virtue, everything on which human happiness on the largest scale depends." Given that lying always weakens this transcendently important ability to some degree, lying will only very rarely produce better outcomes than truth-telling.<sup>43</sup>

Mill of course allowed that the rule against lying would have exceptions, notably for cases in which "the withholding of some fact (as of information from a malefactor, or of bad news from a person dangerously ill) would preserve some one (especially a person other than oneself) from great and unmerited evil, and when the withholding can only be effected by denial." The last clause is interesting, since it suggests that the exception for lying only kicks in if the withholding of information cannot be effected without denying what one takes to be the truth—i.e., without lying outright. One way of withholding information without denying it is to untruthfully implicate. For instance, in the above example of the deathbed question, the respondent withholds the information that the woman's son has died without denying that he has died (as the respondent would if she said, "He is happy and healthy."). So it looks as though Mill's exception to the rule of thumb against lying suggests another rule of thumb: when some rule of thumb tells you to convey something you believe to be false, then, other things being equal, untruthfully implicate rather than lie.

How might we justify a rule of thumb like this? Keeping focus on the damage that either lying or untruthfully implicating does to our ability to trust one another, and to "the cultivation in ourselves of a sensitive feeling on the subject of veracity,"<sup>44</sup> we might argue along the following lines.

One difference between lying and untruthfully implicating is that it is generally easier to be certain that someone has lied than that they have untruthfully implicated, since with implicating there may be room for uncertainty about what has been implicated, whereas there is not typically much room for uncertainty about what has been asserted. As Adler points out, this difficulty in determining and ascribing responsibility should not be confused with a lessening of responsibility.<sup>45</sup> But from the present Millian point of view, it may be relevant that lies are more easily identified than untruthful implicatures and thus are more likely to negatively affect our trust in each other. This might be a reason to make it a rule to untruthfully implicate rather than lie in situations where another rule

---

43. Mill (1863: 32-3).

44. Mill (1863: 32).

45. Adler (1997: 446).

tells you to convey something you believe to be false.

Those who untruthfully implicate also have an easier time denying their dishonesty, not only to others but also to themselves. Indeed, it seems plausible that many choices of untruthfully implicating over lying are motivated by a desire to be able to consider oneself an honest person. Moreover, by refraining from explicitly expressing something one believes to be false, one likely avoids some immediate feelings of shame. None of this makes the choice to untruthfully implicate honorable, but it may well make it beneficial in cultivating our sensitivity about truth telling. By untruthfully implicating instead of lying, one reminds oneself that one cares about truth and is an honest person. In avoiding the more intense shameful feelings produced by outright lying, one protects oneself from experiencing oneself so vividly as morally deficient. These effects may promote one's continued efforts to be honest, whereas choosing to lie might harden one to deception and dishonesty.

Relatedly, it may be that outright lies, as compared with untruthful implicatures, make those lied to feel worse if they discover the lies. Both Adler and Saul point out that, at least in many cases, the effort to tell the truth to someone, even while trying to deceive her, shows greater respect for her than would just lying to her. At least, it usually shows that one is uncomfortable with what one is doing in attempting to deceive, that one is "squirming," as Adler puts it.<sup>46</sup> When the truth is discovered, it may be that the deceived's memory of this effort will make her feel that the gravity of what was done to her was taken seriously. This might help to preserve her trust in others despite having been deceived.

Of course, all of this relies on the truth of the suggested empirical generalizations about the effects of lying versus untruthfully implicating on our sensitivity to truthfulness and trust in each other. (It also depends on a defense of these things as contributing immensely to the goodness of outcomes overall, as Mill claimed they do.) I have offered some reasons for thinking that the correct empirical generalizations might favor untruthfully implicating. But one might argue in the other direction, and not unreasonably, it seems to me.

First, the easier identifiability of lies compared with untruthful implicatures might contribute positively to our trust in one another by making it easier to identify and discuss instances of dishonesty. Second, instead of hardening people to deception and dishonesty, the undeniability of the lie and the shame associated with it might force people to acknowledge their dishonest moments and aim to make them as few as possible. By contrast,

---

46. Saul (2012: 78), Adler (1997: 452).

the ease of self-deception associated with untruthfully implicating might lead to complacency and shallowness about the importance of being honest and forthcoming. And finally, it may be that those on the receiving end of untruthful implicatures have their trust in others diminished *more* than those who are lied to. They may become more inclined to see others as crafty cheaters who want to have their cake by deceiving and eat it, too, by maintaining deniability.

My hunch is that the empirical facts favor the first line of argument rather than the one in the opposite direction. But this is only a hunch. Quite likely, there is extant research in psychology and behavioral economics that would provide a useful basis from which to investigate the plausibility of either argument, or to develop a different one that would support or undermine a rule of thumb like the one Mill's comments suggest. I cannot pursue this here, but it seems like a promising route to justifying one form of the Difference Intuition (or the Reverse Difference Intuition) in a different way than has so far been attempted.

## 6. Conclusion

I began this chapter by trying to clarify the Difference Intuition: the somewhat widespread, somewhat firmly held sense that there is some moral advantage to untruthfully implicating over lying. I defined lying in terms of asserting, but remained open about what precise definition best captures our ordinary notion. Following Grice, I defined implicating as a way of meaning something *without* asserting it. I narrowed down the kind of untruthful implicating that should be compared with lying for purposes of evaluating whether there is a moral difference between them. Just as lying requires a robust form of assertion, so the kind of untruthful implicating to be compared with lying requires a robust form of implicating. I set out a number of ways in which the Difference Intuition might be sharpened. Most of these involve quantifying over contrast cases differing only with respect to lying versus untruthfully implicating, claiming that for all, most, or some such pairs, the case of lying is morally worse than the case of untruthfully implicating. But I also noted that the Difference Intuition could be sharpened to a claim about how we should make moral decisions, rather than a claim about contrast cases, and at the end of the chapter sketched a Mill-inspired strategy for justifying the Difference Intuition in this form.

Before that, I surveyed three broad approaches to arguing that the lying/untruthfully

implicating difference can make a difference to an act's moral status (whether in all cases, most cases, or some cases). The upshot of this survey is, I think, a cautious pessimism. Previous attempts come up short, and the additional possibilities that I suggested do not look especially promising. This, of course, does not show that no justificatory project along these lines could succeed. It only gives a flavor for the kinds of problems that need to be avoided. I am somewhat less pessimistic about the Millian approach outlined in section 5. While this approach would need much more development than I have given it, it has the advantage of taking account of whatever empirical bases we may have, now or in the future, for thinking that some ways of speaking against our minds are in general more detrimental to the overall good than others. It may well be our sense of such generalizations that drive the intuition of moral difference in the first place.

## References

- Adler, J. (1997) Lying, Deceiving, or Falsely Implicating. *Journal of Philosophy* 94(9): 435-52.
- Benton, M. (forthcoming). Lying, Belief, and Knowledge. In Jörg Meibauer (ed.), *The Oxford Handbook of Lying*. Oxford University Press.
- Brandom, R. (1994). *Making It Explicit* (Cambridge, MA: Harvard University Press).
- Carson, T. (2006). The definition of lying. *Noûs* 40 (2):284–306.
- Chisholm, R. & Feehan, T. (1977). The Intent to Deceive. *Journal of Philosophy* 74(3): 143-59.
- Cuneo, T. (2014). *Speech and Morality: On the Metaethical Implications of Speaking*. Oxford: Oxford University Press.
- Fallis, D. (2009). What Is Lying? *Journal of Philosophy* 106 (1):29-56.
- Goldberg, S. (2015) *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.
- Green, M. (2007). *Self-Expression*. Oxford: Oxford University Press.
- Green S., (2006). *Lying, Cheating, and Stealing: A Moral Theory of White-Collar Crime*. Oxford: Oxford University Press.
- Grice, H. P. (1957). Meaning. *Philosophical Review* 66 (3):377-388.
- Grice, P. (1989). *Studies in the Ways of Words*. Cambridge, USA: Harvard University Press.
- Kagan, S. (1988). The Additive Fallacy. *Ethics*, vol. 99 (1): 5-31.
- Keiser, J. (2016). Bald-faced lies: how to make a move in a language game without making a move in a conversation. *Philosophical Studies* 173 (2):461-477.
- Korsgaard, C. (1986). The right to lie: Kant on dealing with evil. *Philosophy and Public Affairs* 15 (4):325-349.
- MacAulay, M. and Stainton, R. (2013), review of Saul (2012) in *Philosophy in Review*, XXXIII(5): 403-5.
- Macintyre, A. (1995) Truthfulness, Lies, and Moral Philosophers. In G. Peterson, ed. *The Tanner Lectures on Human Values*. Volume 16. Salt Lake City: University of Utah

Press.

- Mahon, James Edwin, "The Definition of Lying and Deception", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/lying-definition/>>.
- Mill, J.S. (1863). *Utilitarianism*. London: Parker, Son and Bourne, West Strand.
- Rees, C. (2014). Better lie! *Analysis* 74 (1):69-74.
- Rescorla, M. (2009). "Assertion and its Constitutive Norms," *Philosophy and Phenomenological Research* 79(1): 98–130.
- Saul, J. (2012). *Lying, Misleading, and What is Said*. Oxford: Oxford University Press.
- Searle, J. (1969). *Speech Acts* (Cambridge: Cambridge University Press).
- Shiffrin, S. (2016). Bedrock (Responding to Leslie Kendrick, How Much does Speech Matter?). *Harvard Law Review Forum*: 242-57.
- Shiffrin, S. (2014). *Speech Matters: On Lying, Morality, and the Law*. Princeton and Oxford: Princeton University Press.
- Slote, M. (2007). *The Ethics of Care and Empathy*. London and New York: Routledge.
- Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive by. *Pacific Philosophical Quarterly* 88 (2):251-264.
- Stokke, A. (forthcoming). Lies, Harm, and Practical Interests. *Philosophy and Phenomenological Research*.
- Stokke, A. (2013) Saying too Little and Saying too Much: Critical Notice of Lying, Misleading, and What is Said, by Jennifer Saul. *Disputatio* (35):81-91.
- Stokke, A. (2013). Lying and Asserting. *Journal of Philosophy*, CX(1), 33–60.
- Strudler, A. (2010) The Distinctive Wrong in Lying. *Ethical Theory and Moral Practice* 13(2): 171-9.
- Turmunkh, U., van den Assem, M. & van Dolder, D. (2017) Malleable Lies: Communication and Cooperation in a High Stakes TV Game Show. Available at SSRN: <https://ssrn.com/abstract=2919331>.
- Webber, J. (2013). Liar! *Analysis* 73 (4):651-659.
- Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton: New Jersey: Princeton University Press.