

Lying, Deception, and Epistemic Advantage¹

Eliot Michaelson
King's College London

Andreas Stokke
Swedish Collegium for Advanced Study
Uppsala University²

To appear in
Khoo and Sterken (eds.), *The Routledge Handbook of
Social and Political Philosophy of Language*,
Routledge, forthcoming

The Bishop. Can Alexander explain why one lies?
Alexander. One lies to win an advantage.
- I. Bergman, *Fanny and Alexander*

1 Introduction

Lying is a paradigmatic way of deceiving with words. But there are other ways of using language deceptively that don't amount to lies. Bernard Williams gives the following example:

“Someone has been opening your mail,” she helpfully says, and you, trusting her, take it that it was not the speaker herself. If you discover that it was the speaker, you will have to agree (if through clenched teeth)

¹ We are grateful to Don Fallis, Ethan Nowak, Quinn White, Torfinn Huvenes, and Vlad Krstić and for helpful comments and discussion.

² This work is entirely collaborative and names appear in alphabetical order.

that what she said was true. So, you must also agree, she did not tell you a lie. (Williams, 2002, 96)

With respect to Williams's example, the contrast between lying and speaking deceptively without lying – often referred to as 'merely misleading' – might be summarized as the difference between uttering (1a) or (1b).

- (1) a. I haven't been opening your mail.
- b. Someone has been opening your mail.

In this situation, uttering (1a) is to lie while uttering (1b) is not. Crucially, (1a) is something the speaker believes (indeed knows) to be false, whereas (1b) is something she believes to be true. Yet both utterances are aimed at the same thing: deceiving the hearer into believing that the speaker has not been opening the mail.

Based on this sort of contrast, philosophers throughout history have commonly held that to lie is to say something one believes to be false with the intent to deceive one's listeners into believing it. You lie only if you say something you believe to be false, as in (1a), and as opposed to (1b). But moreover, according to this tradition, the liar is centrally concerned with instilling a false belief in the hearer. So, on this broad view, lying is a form of deception that is characterized by explicitly saying what one wants to deceive one's listener about.

This conception of lying can be further justified via consideration of a wide range of commonplace lies. Most lying fits the bill of saying something one believes to be false with the intention to make one's addressee believe it. In an attempt to impress your friends, you tell them that you went camel-riding during your summer holidays while knowing full well that you stood by and watched. You tell your housemate that the market was out of butter, even though you simply forgot to buy it. Or perhaps you find yourself telling your colleague that you enjoyed reading his or her recent book, even though you did not care for it.

Unsurprisingly, given all this, the intention to deceive has long been viewed as a central – perhaps even *the* central – characteristic of lying. In the 4th century Augustine wrote,

In reality, the fault of a person who tells a lie consists in his desire to deceive in expressing his thought. (Augustine, (1952 [395]), 55-56)

Many, probably even most, philosophers since have followed Augustine in understanding the intention to deceive as a central – perhaps even *the* central – moral problem with lying.

In the last two decades, however, this philosophical package has been called into question. In light of examples involving what are commonly known as ‘bald-faced lies’, many have found the traditional view wanting. That has served to put the question of what it is that makes an utterance a lie at the center of recent debates regarding the morality of lying.

This chapter first rehearses this dialectic, including some recent responses on behalf of traditionalists. We argue that none of these responses look very promising. Second, we argue that neo-traditionalist attempts to save the traditionalist account of lying ultimately threaten to hollow out the notion of deception, divorcing it from its most interesting features.

By our reckoning, the central feature of an intention to deceive is that it aims to give the deceiver an *epistemic advantage* over the deceived. There is, however, no way for neo-traditionalists to endorse this claim about deception. That, in turn, threatens to undermine one of the main advantages of the view: its purported ability to explain the wrongness of lying. We close with what we hope will be a note of reconciliation: whether or not one accepts the Augustinian picture, one can accept that one of the central motivations for lying is to gain an advantage over one’s audience. This raises the possibility that what makes lying wrong, when it is wrong, is that this advantage is sought out unjustly.

2 Bald-Faced Lies

2.1 The Augustinian Definition

Augustine’s work on lying spawned a tradition which defined lying as a species of deception. More precisely this view can be spelled out as follows:

The Augustinian Definition of Lying

A lies to B if and only if there is a proposition p such that:

(LD1) A states that p to B, and

(LD2) A believes that p is false, and

(LD3) By saying that p to B, A intends to deceive B into believing that p .

Against this traditional view, a number of philosophers have recently argued that lying does not necessarily involve intentions to deceive.³ That is, even if you do not aim at deceiving your listener, you might still be lying. One central

³ E.g. Carson (2006), (2010) Sorensen (2007), Fallis (2009), Stokke (2013).

type of counterexample to the traditional view involves what Roy Sorensen (2007) has called 'bald-faced lies.'

Here is (a version of) a widely discussed example originally introduced by Thomas Carson (2006, 290):⁴

The Cheating Student

The Dean is terrified of lawsuits. Accordingly, she will only punish students for cheating if those students admit to having cheated. The Dean's policy is well-known to all parties. Professor X has impeccable video evidence of a student cheating during the exam. The student, Professor X, and the Dean have all just watched this video, and are all aware that the others watched as well. So it is common knowledge that the student cheated on the exam. The Dean asks the student if she cheated on the exam and the student replies "I did not cheat on the exam."

Carson points out that, in cases like this one, condition (LD3) of the Augustinian Definition is not satisfied. Since it is common knowledge that the student cheated on the exam, it would seem that the student cannot intend to deceive the Dean into believing that she did not cheat on the exam. Not only does the student know that the Dean knows that she cheated, she also knows that the Dean knows that she knows that the Dean knows that she cheated. So how could she hope to convince the Dean otherwise?

If this line of reasoning is right, (LD3) fails and the Augustinian Definition predicts that the student was not lying. And yet, most take it to be abundantly clear that she *was* lying.

Two main strategies have emerged for defending the Augustinian Definition of Lying in response to cases like the Cheating Student. Some writers have responded by denying that putative bald-faced lies, like the student's utterance, really are lies. Others have accepted that bald-faced lies are *bona fide* lies while arguing that bald-faced lies are deceptive, initial appearances to the contrary

In the rest of this section we consider the first of these strategies. We devote Section 3 to the second of them.

2.2 Bald-Faced Lies and Assertion

There are different ways of arguing that bald-faced lies are not genuine lies.⁵ Here we focus on a particular type of argument according to which utterances

⁴ Carson (2010, 21) offers a slightly different version of the example.

⁵ See e.g. Mahon (2011), Meibauer (2011), (2014), Faulkner (2013), Keiser (2015).

like the student's utterance to the Dean are not *assertions*. There is broad agreement in the literature that you lie only if you make an assertion—though not everyone uses that particular terminology. According to this line of thought, since bald-faced lies are not assertions, they are not lies. Therefore such examples do not show that you can lie without intending to deceive.

One example of this strategy can be found in work by Jörg Meibauer (2014). His main argument for the claim that bald-faced lies are not assertions is that the bald-faced liar

does not really present p as true in the context since he *lets shine through* that p is false. He would not feel committed to the truth of p , and he would not be ready to provide further evidence. (Meibauer, 2014, 140)

In other words, Meibauer proposes that (LD1) should be replaced with something stronger, something like: A assertively commits herself to B regarding the truth of p . The cheating student does not assertively commit herself to the truth of p in this way—so she isn't lying. Hence, the case fails to constitute a counterexample to the revised Augustinian Definition.⁶

The problem with this line of argument is that it subtly shifts the original example. Even though the student knows that the Dean knows that what she is saying is false, the student in no way signals this mutual knowledge in making her utterance. So there is no reason to think that the student “lets shine through” that her statement is false. This stands in contrast to uttering something while winking or speaking in an ironic tone of voice. Moreover, the student here is arguably committed to her statement in the sense that, subsequently, she is in no position to claim that it wasn't uttered in earnest. And, finally, while the student may not be in a position to provide further evidence to support her claim, because there is none, she would surely do so if she could.

A more promising defense of the claim that utterances like the student's are not assertions, and hence are not lies, is offered by Ishani Maitra (2018). Her argument is founded on a view of assertion according to which assertion is a speech act governed by the following constitutive rule:

Evidence-Responsiveness Rule: If a speaker S 's utterance of U is not sufficiently responsive to her (total) evidence that bears on p , she does not assert p via uttering U . (Maitra, 2018, 72)

Most philosophers agree that, as Timothy Williamson puts it,

⁶ Cf. Stokke (2018, ch. 1).

Constitutive rules do not lay down necessary conditions for performing the constituted act. When one breaks a rule of a game, one does not thereby cease to be playing that game. (Williamson, 1996, 491)

Similarly, on Maitra's view, one can fail to conform to the Evidence-Responsiveness Rule (henceforth, ERR) and nevertheless succeed in making an assertion. Yet, according to Maitra, some violations of ERR do thereby constitute failures to make an assertion. In particular, such failures must be *flagrant*, in a particular sense. Hence, on Maitra's account,

flagrant failures to conform to this rule are *incompatible* with asserting, while both conformity to the rule, as well as *non-flagrant* failures to conform, are compatible with asserting. (Maitra, 2018, 72)

So, flagrant violations of ERR entail that the speaker is not asserting at all, but mere violations of ERR, in contrast, do not.

Having clarified this, we can now see how Maitra will explain why bald-faced lies like that of Carson's cheating student constitute merely apparent counterexamples to the Augustinian Definition. Namely, the cheating student is not making an assertion. Hence, *pace* our initial reaction to the case, it should not ultimately be viewed as a counterexample; we just had to get clearer on what it takes to assert something in a context.

Supposing for the moment that ERR is well-motivated, the persuasiveness of this response clearly depends on how convincingly one can make the case that apparent bald-faced lies involve flagrant violations of ERR. So we should ask: what makes a violation of ERR flagrant?

Maitra repeatedly says that the failure must be "intentional and sufficiently marked." (Maitra, 2018, 72 *et passim*) With respect to bald-faced liars in particular, Maitra claims that, rather than make assertions, "they do something much more like what an actor does." (Maitra, 2018, 65) By way of motivating this claim, Maitra mentions four points of comparison between the cheating student and an actor on stage:

[(i)] The student knows that his utterance ("I didn't cheat on the exam") is false. Moreover, he also knows that it is not at all supported by his evidence. After all, he knows that he cheated. The Dean knows these things too. (Maitra, 2018, 66)

[(ii)] [C]ome what may, he will continue to insist that he didn't cheat on the exam. (Maitra, 2018, 76)

[(iii)] [T]he student's performance is only partly intended for his immediate interlocutor, the Dean. Crucially, it's also intended for others beyond his immediate audience, such as future readers of any transcript or summary of the proceedings. (Maitra, 2018, 77)

[(iv)] Fourth, and finally, each of the student and the Dean is aware that the other is aware of the features of the student's performance enumerated above. (Maitra, 2018, 77)

We note that, of these, only (iv) and the observation that "The Dean knows these things too" in (i) are distinctive of the bald-faced liar as opposed to ordinary, deceptive liars. Uncontroversially, ordinary deceptive liars know (or at least believe) that what they say is false and is not supported by their evidence. Likewise, ordinary liars will typically continue to insist on the truth of what they say. But moreover, there are arguably many instances of ordinary, deceptive lying that are intended for an audience beyond the immediate interlocutor—such as, politicians lying in meetings attended by reporters, and so on.

In other words, Maitra's argument that bald-faced lies are not assertions rests on the claim that the bald-faced liar and her audience are mutually aware that the speaker is not responding to evidence, in Maitra's sense. This is what makes the violation of ERR flagrant, and hence what makes the utterance fall short of assertion.

However, even if one is sympathetic to a view of assertion as governed by one or more constitutive rules, there are reasons to doubt the suggestion that asserting is incompatible with flagrantly violating such a rule, in this sense. That is, there are reasons to reject the idea that if the speaker and hearer are mutually aware that the speaker is violating the relevant constitutive rule, the utterance no longer counts as an assertion. In general, given that constitutive rules ought not be understood as necessary conditions for playing the relevant games, it is unclear why instances of breaking such rules with the added feature of mutual awareness would amount to opting out of the game.

It is a constitutive rule of soccer that deliberately handling the ball (by any player who is not goalkeeper) is a foul. Consider the following story:⁷

Flagrant Handling

Lisa is deemed to be offside during an attack. She is upset and angry with the referee who she thinks made a mistake. In her rage she looks

⁷ Thanks to Torfinn Huvenes for this example.

the referee straight in the eye and in plain view of everyone on the field and all the spectators she picks up the ball and throws it away.

Lisa is clearly violating the no-handling rule, and she is clearly doing so flagrantly, in Maitra's sense. But it does not seem to us that she is thereby ceasing to play the game. What will happen as a result of Lisa's violation of the rule? Most likely, she will get a yellow card and a free kick will be awarded the opposing team. But the game has not stopped. She has not placed herself outside the rules by breaking the no-handling rule flagrantly.

This (not unrealistic) situation is significantly different from situations that would be clear instances of opting out of the game. Suppose, for instance, that instead of picking up the ball, Lisa in her rage just walks off the field, and out of the stadium, refusing to listen to anyone or take any further part in the playing of the match. Or imagine that she sits down on the field, refusing to get up or move or do anything. These are clear cases of opting out of the game. Importantly though, neither of them involves violating any constitutive rules. Rather, opting out of a game is most commonly a matter of no longer performing actions that typically constitute playing it.

Similarly, if one thinks that assertion is governed by one or more constitutive rules, and one moreover thinks that liars like the cheating student, along with ordinary liars, are violating such rules, one should not think that doing so flagrantly is tantamount to not making an assertion. There are any number of ways of uttering a declarative without making an assertion—winking, for instance, or impersonating someone else—but, as best we can tell, stating something that is not responsive to one's evidence is not one of these ways, even when this non-responsiveness is apparent to everyone involved, and known to be so. Put slightly differently: if one isn't asserting p in saying it, then it won't be appropriate to hold one responsible for putting forward that content in the ordinary way. But we can see no reason to think that one's failing to be properly responsive to the evidence—even when that is mutually evident to all involved—should prove exculpatory in this manner. On the contrary, these strike us as some of the most important instances in which to be able to hold each other to account for asserting this or that.

3. Lying and Intentions to Deceive

We noted two main strategies for defending the Augustinian Definition of Lying in the face of bald-faced lies. In the previous section we considered a version of the first of these, centering on the claim that bald-faced lies are not

assertions and hence are not lies. We now turn to the second general strategy, according to which bald-faced lies are indeed genuine lies, but should nevertheless be seen as deceptive in some way.

3.1 Deception and Deceit

The main proponent of this line of argument is Jennifer Lackey (2013), who attempts to defend a neo-traditionalist picture of lying by substituting a broader notion of 'deception' for the more traditional 'deceit.' She offers the following sufficiency claims for each of these notions:

Deceit: A deceives B with respect to whether p if and only if A aims to bring about a false belief in B regarding whether p .

Deception: A is deceptive to B with respect to whether p if A aims to conceal information from B regarding whether p . (Lackey, 2013, 241)

The former notion of deceit is the one involved in the Augustinian Definition, more precisely in (LD3). So, since Lackey agrees that bald-faced lies are lies, she rejects (LD3) and proposes that lying instead be characterized in terms of her notion of deception.

Here is Lackey's (2013, 237) definition of lying:

Lackey's Definition of Lying

A lies to B if and only if

(LL1) A states that p to B, and

(LL2) A believes that p is false, and

(LL3) A intends to be deceptive to B in stating that p .

Lackey's condition (LL3) contrasts with the traditional (LD3), according to which lying necessarily involves the intention to induce false beliefs in the audience. Hence, on Lackey's view, lying requires saying something you believe to be false and thereby intending to be deceptive to your listener. But it does not require intending to deceive her.

Returning to the case of the cheating student, Lackey's key claim is that, even if the student cannot intend to deceive the Dean in the sense of making the Dean believe that she did not cheat, she can still intend to conceal information from the Dean regarding whether she did. More specifically, the student can intend to conceal her confession that she cheated on the exam. Assuming that such a confession, were there one, would count as evidence, and assuming further that evidence counts as information, Lackey's neo-traditionalist

definition of lying will classify this case as a lie. The student intends to conceal her confession, and thus (LL3) is fulfilled.

3.3 Concealment

Don Fallis (2015) and Andreas Stokke (2018) have each criticized Lackey's position by arguing that the cheating student cannot be said to aim at concealing information from the Dean, and hence does not satisfy (LL3). Lackey specifies that

Concealing information regarding whether p can be understood broadly here, so that it subsumes, among other phenomena, concealing *evidence* regarding whether p . (Lackey, 2013, 241)

Further, Lackey argues that

while he does not intend to deceive the Dean into falsely believing that he did not cheat, he does intend to conceal crucial evidence from the Dean that is needed for punishment from the university—namely, an admission of wrongdoing. (Lackey, 2013, 241-242)

Yet, as Fallis (2015) argues, this argument relies on the questionable assumption that one can conceal x even if x does not exist:

[T]he student does not aim to *conceal* his confession. There is no confession to be concealed since he has not confessed. It is not like when a criminal conceals or destroys existing evidence of his crime. (Fallis, 2015, 90)

Similarly, if Jones does not have, or have access to, the blood-stained dagger, she cannot conceal it from the police. And, even more clearly, if there is no blood-stained dagger at all, no one can conceal it, or even intend to do so, as long as they are aware of its non-existence.

Many will conclude that the cheating student cannot be said to intend to conceal her confession from the Dean, since there is no confession to conceal. So to defend the view that the student is deceptive, in Lackey's sense of (LL3), one might instead try to argue that the student intends to conceal information from the Dean in some other sense.

There are two main candidates for which information the student might be said to intend to conceal:

- (2) a. That the student cheated.

b. That the student believes she cheated.

As suggested earlier, however, there are plausible versions of Carson's original example on which both (2a) and (2b) are common knowledge. But if both these are common knowledge, then it is unclear how the speaker might intend to conceal either of them. In general, it is plausible to think that if p is common knowledge between A and B, then, since A knows that B already knows p , A cannot rationally believe that she can conceal p from B. If Jones knows that the police know that she knows that they know (etc.) that she has the dagger, she cannot rationally believe that she can conceal the fact that she has the dagger from the police. From this, Fallis and Stokke both conclude that, assuming that the student is a rational agent, she cannot intend to conceal either (2a) or (2b) from the Dean.

This argument relies on the assumption that, roughly, one cannot intend to do what one believes one cannot do. That is, on some version of the following widely endorsed principle:⁸

Belief-Intention Constraint

A intends to φ only if A believes that she can φ .

At this point, if one wants to defend the kind of neo-traditionalist view embodied by (LL3), there are two main options. First, one could opt to broaden the notion of deception further still, so that the student counts as being deceptive after all. Second, one could opt to broaden the notion of intending beyond the Belief-Intention Constraint, so that the student counts as intending to conceal information in spite of her believing that she will inevitably fail. We will consider each in turn.

3.4 Withholding Information and Deception *Secundum Quid*

One way of modifying the neo-traditionalist notion of deception is to focus on withholding information instead of concealment. This might be spelled out as follows:

Deception*: A is deceptive to B with respect to whether p if A aims to withhold information from B regarding whether p .

⁸ See e.g. Audi (1973), Grice (1973), (1989, 98), Davidson (1989 [1978]), Velleman (1989), Neale (2005), Mele (2010), Stokke (2010), (2018), Kissine (2013, ch. 2), Fallis (2015). See also Velleman (1989, 113–115) for discussion and more references.

The first thing to note about this suggestion is that, in order to succeed, it would need to be argued that causing false beliefs is a form of withholding information. If not, the neo-traditionalist view will fail to capture the sort of ordinary deceptive lies that originally motivated the Augustinian view. In other words, if the neo-traditionalist wants to appeal to Deception* in fleshing out her claim that lying necessarily involves an intention to be deceptive, she needs to argue that whenever A aims to cause a false belief in B, A necessarily aims to withhold information from B.

To make good on this challenge, it might be suggested that when causing someone to falsely believe that *p*, one is thereby preventing them from acquiring the true belief that not-*p*. Yet to argue that all cases of causing false beliefs are cases of preventing acquisition of true beliefs is arguably a tall order. Within their category of commissive deception, Chisholm and Feehan (1977) distinguished between what they called positive deception *simpliciter* and negative deception *secundum quid*. We spell out this distinction as follows:

Positive commissive deception *simpliciter*: A contributes causally to B's acquiring the false belief that *p*.

Negative commissive deception *secundum quid*: A contributes causally to preventing B from acquiring the true belief that not-*p*.

At this point, then, the neo-traditionalist is forced to argue that the former entails the latter. Unsurprisingly, though, there are counterexamples to this proposal.

Consider the following elaboration on Williams's example involving (1a):

Mail Opener

Anna has been opening Sheryl's private mail while she's been away. Sheryl's back in the office, and she sees that someone has been opening her mail. She asks Anna who it was. As it happens, Anna has just quit her job in the office that very morning, and she is on her way to leave for good and will, in all likelihood, never see Sheryl or anyone else from the firm again. She also knows that their colleague, Emily, will appear any second now and reveal the truth to Sheryl. So, Anna says, "I haven't been opening your mail." Sheryl thanks her and wishes her well in her future career, and Anna immediately leaves the office. A short while later Emily reveals the truth to Sheryl, as Anna predicted.

Anna's utterance is clearly a lie. Moreover, her aim is clearly to make Sheryl believe the falsehood that she has not been opening her mail so that she can leave the office quickly. But she is clearly not preventing Sheryl from learning the truth. And she knows full well that she is not doing so. Hence, as *per* the Belief-Intention Constraint, she cannot be said to intend to contribute casually to preventing Sheryl from acquiring the relevant true belief.

There is more room to maneuver here. For instance, the neo-traditionalist might find some way of making good on the claim that intending to cause a false belief is invariably a way of intending to withhold information. Yet, as we will see next, there are further problems that arise if one opts for a definition that runs off of Deception*, either in total or just in part.

3.5 Minimal Deception and Epistemic Transparency

Even if a conception of deception along the lines of withholding information can be made to capture the traditional causal category, the question remains whether bald-faced liars can be said to have intentions to withhold information. Lackey explicitly glosses withholding information as distinct from concealment:

To withhold information is to fail to provide it, rather than to hide or keep it secret. (Lackey, 2013, 241)

Fallis argues that the cheating student does not fit into this category:

But the student does not aim to conceal, or withhold, any information that is relevant to his guilt. If an object (or an event) might convey some information about some topic, but you already have that information, I cannot withhold the information from you by withholding the object (or the event). (Fallis, 2015, 90)

This is right only to the extent that withholding information necessarily involves preventing someone from acquiring it. By contrast, as quoted above, Lackey suggests that to withhold information is "to fail to provide it."

There is arguably a sense in which the student fails to provide the relevant information. In particular, she fails to provide the information in (2a), i.e. that she cheated. To be sure, she is not preventing the Dean from acquiring (2a). The Dean already has that information; indeed, the Dean knows (2a). Rather, by failing to provide it, she is interfering with the Dean's ability to make use of the information they both already have. Hence, there is arguably a sense in which the cheating student does withhold information from the Dean (and intends to do so). So, if the neo-traditionalist wants to construe the kind of

deception involved in lying as Deception*, it is plausible to think that (at least many) bald-faced liars can be characterized as deceptive, after all.

Yet this only vindicates the neo-traditionalist claim to reinstate the traditional view of lying as a species of deception to the extent that Deception* is a plausible notion of deception. As we will now argue, there are reasons to be skeptical about this.

If the cheating student can be seen as deceptive even in the case where the relevant information is common knowledge between her and the Dean, this means that you can intend to deceive someone without intending to tamper with their epistemic state regarding the information in question. If this strikes one as implausible, it might be because one finds the following more general idea attractive:

Minimal Deception

In ϕ 'ing, A is deceptive to B regarding p only if, in ϕ 'ing, either A contributes causally to B being wrong about p or A allows B to continue being wrong about p .

By being wrong about p we mean, roughly, having bad beliefs or credences regarding p . In terms of full belief, this means having false beliefs. In terms of credences, roughly, we have in mind cases in which one's actual credences do not match probability on one's evidence, or rational credence. For instance, if p is .6 probable on one's evidence, having a credence of .4 or .8 are both ways of being wrong about p , on this view. If one does not intend to be minimally deceptive in this sense, then one does not have an intention to deceive.⁹ This is arguably a way of capturing the general thought behind Chisholm and Feehan's taxonomy of deception and deceptive intentions.

To deny Minimal Deception is to accept that deception can be epistemically transparent. That is, on such a view, A can deceive B simply by telling B something she thinks is false, even though they both know that they both know (etc.) that it is false, and even though A knows full-well that telling B the falsehood will have no effect whatsoever on B's beliefs or credences regarding what she says. This, it seems to us, is to effectively abandon the concept of deception. (We return to some further ramifications of this thought in Section 4.)

3.6 Trying

⁹ Note that one can intend to be minimally deceptive with respect to a content other than the one asserted—for instance, an implicature.

We noted that there are two main ways of defending a neo-traditional view based on (LL3) in order to argue that bald-faced lies involve intentions to deceive. We have seen that a challenge for the first of these is to avoid giving up Minimal Deception. We now turn to the second strategy, that of modifying the notion of intending.

The arguments considered so far have relied on the Belief-Intention Constraint. The general thought is that, since (2a-b) are common knowledge between the student and the Dean, the student cannot (rationally) believe that she can conceal them from the Dean. Hence, given the Belief-Intention Constraint, she cannot intend to do so, and therefore the student does not satisfy the neo-traditionalist proposal as embodied in Lackey's notion of Deception.

One strategy, then, would be for the neo-traditionalist to reject the Belief-Intention Constraint. This, however, comes at a substantial cost for the view. Effectively, one has saved the view that lying necessarily involves intentions to deceive by rejecting a central tenet of most theories of intentional action.

A more attractive route would be to look for another conative category that is not subject to a constraint like the Belief-Intention Constraint. Consider, for instance, the following modified version of (LL3):

(LL3*) A tries to be deceptive to B with respect to whether p by stating that p .

For the sake of argument, we assume that trying is not constrained by beliefs in the way that intentions plausibly are. That is, we assume that someone can try to φ even if she does not believe she can φ , and hence does not intend to φ .¹⁰ So, the claim would be that, in lying, the cheating student tries to conceal information from the Dean, even if she does not intend to do so. However, as we explain below, there are lies that will incorrectly be ruled out by (LL3*).

3.7 Self-Realizing Lies

Consider a new example:¹¹

Inbox

¹⁰ Here, we are following a line of thought on tryings developed in Ludwig (1992).

¹¹ This example bears some resemblance to the "tell-tale" variety of bald-faced lies that have been discussed by Krstić (2019). The difference is that, in the latter cases, there is a feature (typically a behavioral feature of the liar) which makes it common knowledge that the speaker is lying, whereas in cases like Inbox there is a feature that makes it common knowledge that what the speaker says is false.

Betty, the department secretary, is a master at navigating the university's byzantine administration. Graduate students regularly come to her for help with various obscure forms, which she is generally happy to render if asked politely. One day, one of the less competent graduate students, A, stops by with a form that he has been trying to figure out. Rather than asking politely for help, he thrusts the form at Betty and demands "Betty, you need to fill out this form for me!" Betty calmly takes the form and says to A, "I'll put this at the top of my pile." As she says this, she carefully, deliberately, and in full view of A places the form at the bottom of her overflowing inbox.¹²

Betty's utterance is a bald-faced lie. Significantly, it is made false by the very gesture that Betty makes simultaneous with her utterance. Betty's utterance therefore an instance of what we might call a 'self-realizing' lie. Such lies do not satisfy (LL3*).

First, note that Betty's lie does not satisfy (LL3*) with respect to Deception—that is, the notion according to which deception is a matter of concealing information. With her gesture Betty is actively trying to *reveal* that she is not going to put the form at the top of her inbox. She is not trying to conceal that information. Hence, she is not being deceptive.

Second, note that Betty's lie does not satisfy (LL3*) with respect to Deception*—that is, the notion according to which deception is a matter of withholding information. To say that it does is to say that Betty is trying to withhold information regarding whether she will put the form at the top. But again, by gesturing the way she does, Betty is actively trying to reveal that she is not going to do so. Hence, to say that Betty satisfies (LL3*) with respect to Deception* is to say that Betty's mental state (or perhaps her act itself) is incoherent. However, it is clearly implausible to say that Betty's mental state (or her act) are incoherent. Both her mental state and the act whereby she expresses that state are perfectly comprehensible.

In other words, lies like Betty's do not satisfy (LL3*). By her utterance, Betty is not even trying to be either Deceptive or Deceptive*.

One potential response here would be to argue that Betty's utterance is an instance of irony. In that case, her words would serve to communicate something other than their literal meaning—presumably something true—and thus she would not be lying.

¹² One of the authors would like to take the opportunity to express his profound thanks to Betty White for all of her generous help navigating UCLA's generally Kafkaesque forms and regulations during his time as a PhD student there.

We are skeptical about such a response, however. Whether one is speaking ironically is typically understood as a matter to be decided by appeal to the speaker's mental state. Speakers speak ironically when they have certain sorts of intentions—typically, intentions for their literal meaning to be disregarded in favor of some other communicated content. We are free to simply stipulate that the version of the case in which we are interested is one where Betty lacks whatever intentions would be required to speak ironically; rather, she intends to speak plainly and for her utterance to be plainly and obviously false; she wants the utterance to be heard unironically, as she aims to be slightly rude to this graduate student. Even stipulating all this, we can see nothing incoherent about either Betty's mental state or her act.

Cases like Inbox demonstrate that it is possible to lie while simultaneously pointing to or otherwise making evident the very falsehood of what one is saying. Even if the neo-traditionalist could deal with the Cheating Student and other instances of bald-faced lying, it is harder to see a way for her to deal with self-realizing lies like these. The basic problem is that, on any version of the neo-traditionalist theory, self-realizing lies will incorrectly be classified as incoherent. That, however, hardly seems an accurate description of them.

4 Lying, Deception, and Advantage

4.1 Minimal Deception and Epistemic Advantage

Earlier, we suggested that it is a bedrock fact about deception that it cannot be epistemically transparent. One cannot deceive someone without either causing or helping maintain that person's being wrong about the relevant information. We put this in terms of a condition on deception:

Minimal Deception

In ϕ 'ing, A is deceptive to B regarding p only if, in ϕ 'ing, either A contributes causally to B being wrong about p or A allows B to continue being wrong about p .

We want to highlight a corollary to minimal deception:

Epistemic Advantage

If, in ϕ 'ing, A is deceptive to B regarding p , then A either gains or maintains an epistemic advantage over B with regard to p .

A holds an epistemic advantage over B regarding p just in case A's credence in p matches the probability of p on her evidence better than B's credence matches the probability of p on B's evidence, and A's evidence is at least as good as B's. Given what we said earlier, this amounts to, roughly, A being "less wrong" about p .

Of course, deceptive acts will not necessarily succeed in bringing about an epistemic advantage; the deception might not work, or the agent herself may be significantly confused about how the world is. Rather, when one aims to be deceptive, one aims to gain an epistemic advantage over another; one takes oneself to be in a relatively good doxastic state regarding p , and one aims to leverage that to put one's interlocutor in a worse state regarding the very same proposition.

To illustrate, consider the following situation. The probability of p on both A's and B's evidence is .5. A has .2 credence in p and B has .1 credence in p . A tells B that p . As a result, B's credence increases to .2. It might seem that in this case A is deceptive towards B because she tells her something she regards as unlikely (perhaps false) and yet she causes B to be less wrong about p . However, we take this to be a case in which A aimed at being deceptive towards B, and also thereby aimed at gaining an epistemic advantage over B, but failed at both. Had A realized what B's initial credal situation looked like, we would expect, given her aims, for A not to have said p here. Otherwise, A would correctly be described as aiming to make B epistemically better off, which is arguably not a way of being deceptive.

Since we do not think all lies are deceptive, we do not think that all lies aim at epistemic advantage. The lies told by both Betty and the cheating student do not aim at epistemic advantage. Still, we take it that these lies do generate or maintain a sort of positional advantage—just not an epistemic one. What sort of advantages are these, exactly, and what in general can be said about why lies are potentially appealing ways of obtaining those sorts of advantages?

4.2 Positional Advantage

We characterized epistemic advantage as having a perspective on the world that is, in a certain sense, better or more accurate than someone else's. Intentionally bringing it about that one has an epistemic advantage over another, then, should put one in a position to take advantage of that gap. Lying is a potentially effective tool for bringing about such an advantage because it in no way affects one's own epistemic state, and yet it has the potential to influence one's interlocutor's state for the worse—relative to one's own take on things, at least.

But lies do not just have effects on epistemic states. As speech acts, they can affect all manner of other sorts of mental and social states as well. Even if we were to assume that only assertions of declarative sentences can constitute lies, that form of sentence can be used to do a great deal more than just to trade information back and forth: one can use such sentences to make contracts and commitments (“I’ll see you at the airport”), to name things (“I dub thee the Furzey Gorse”), to make believe (“You enter a dimly lit cavern. In the distance, you hear the muffled voices of orcs.”), etc. And in addition to all this, we can of course use declaratives to prompt people to feel all manner of things, like grief, anger, pride, and so forth.

This broadens the range of things about which one might try to gain an advantage by means of a lie. So, for instance, one might try to obtain an emotional advantage over one’s chess opponent by informing her that you slept with her partner a few days prior. Even if you expect your opponent not to believe you, the ruse may well take up some of her mental energy during the match—potentially leaving her play distracted. Epistemic advantage, then, is only one sort of advantage that can be obtained via lying.

Our earlier cases provide us with some further examples. In the case of the cheating student, for instance, lying deprives the Dean of certain potential courses of action, given the Dean’s unwillingness to act in the absence of a confession. This constitutes a certain sort of positional advantage within the bureaucratic procedures of the university. Betty, on the other hand, blocks A from asking for further assistance without outright telling A that she is unwilling to help him. If A tries to follow up with someone else, Betty can find A’s form in her pile and say something to the effect of “Look, it’s right here. I just haven’t had a chance to get to it yet.”

4.3 Political Lies and Nihilistic Advantage

One rather perplexing question in contemporary politics is why politicians lie so much—particularly about things where there is little to no chance of their being believed. One possibility is that they are trying to have other sorts of doxastic effects, like defeating our knowledge on certain topics. Another is that they are trying to have emotional effects on their listeners—one might even go so far as to say ‘manipulate the emotions’ of their listeners. Surely, both of these are part of the story of lying in recent political life.

But there is another sort of advantage that we think is likely more specific to our contemporary political milieu, and one that we take to be worth drawing attention to. Suppose that one belongs to a party which professes to stand for a range of goals: lowering crime, strengthening the middle-class, and ensuring access to high-quality public education and health care. Suppose

further that many of these outcomes overlap with the professed goals of one's political opponents; the differences between the parties (with respect to this range of goals, at least) amounts largely to differences in strategies for obtaining these goods rather than differences over what the relevant goods are. Finally, suppose that there is a large and growing body of research suggesting that one's policy prescriptions for obtaining these goals are bunk. In contrast, your opponent's prescriptions are generally well-supported by the evidence.

Strategically, in order to take or remain in power, what should one do? One could of course shift one's policy proposals. But then these would fail to distinguish you from your opponent. Alternatively, one could shift the goals towards which one is working. But supposing that these are widely-accepted in the society in which you live, that might well prove costly. Or one could lie.

The more standard calculation is that you might manage to convince some of those listening to you that the available evidence is bogus, or does not actually support your opponent. But there is an even more insidious strategy to be played here. For suppose that it is already prevalent in your society to think that politicians are apt to lie. Then, when you disagree with your opponent about some basic facts, your audience is likely to think that one of you must be lying—regardless of whether they have an opinion on who. That, in turn, should prompt them to raise their credence in the proposition that *politicians often, or even always, lie* even more. In fact, it doesn't even matter if you expect that members of the public are going to brand you a liar. So long as that raises their credence in propositions like *politicians are liars*, that serves to hurt your opponent. For one of your opponent's main advantages is that she actually has better policy prescriptions than you do, at least as far as one can tell given the evidence. But if informed debate on that is off the table—if both are lying, then what's the point of trying to figure out the truth?—then the election will have to hinge on other things. Since one was peddling make-believe from the get-go, this is purely to your advantage.

We propose to call this 'nihilistic advantage': a sort of positional advantage that obtains when one's goals are aided by others coming to believe that everyone is lying to them, or perhaps even that there is no truth at all. Sadly, given a certain set of background credences and social conditions that might very well obtain in our own world, one can easily see how there is going to be a fairly direct way of playing for such advantage by lying, and blatantly lying, to the public. When you have nothing but fantasy to sell, the truth—and, in particular, the thought that the truth might be readily obtained by either politicians themselves or at least the media—becomes the enemy. A widespread distrust that truth can be obtained, in turn, becomes a positional advantage. Sometimes, then, one can play for positional advantage by trying to convince one's audience that everyone—even, possibly, oneself—is a liar.

5. Conclusion

We began by considering the traditional, Augustinian account of lying, along with some recent challenges to that account. Next, we considered two broad strategies for responding to those challenges. Finding both of those strategies wanting, we urged caution with respect to neo-traditionalist theories of lying. What most interested us, however, was the worry that, in trying to defend a neo-traditionalist account, one might be forced to give up on the plausible claim that deception aims at generating epistemic advantage. Finally, we suggested that, if we are willing to accept a wider notion of ‘positional advantage’ that subsumes epistemic advantage but is by no means equivalent to it, then we might well derive a fairly unified account of what makes lies wrong: lies are wrong when one uses them, unjustly, to seek out some positional advantage.

So far as we can tell, this is both an attractive and underexplored thesis. This latter fact is somewhat surprising, particularly given how ecumenical this thesis turns out to be; depending on how hard a line we take regarding when one might be justified in seeking positional advantage, the thesis should be equally amenable to those who think it is never justified to lie and those who think it is often justified. What might have seemed like a slugfest of intuitions about lying is revealed to hinge on a deeper disagreement about what might justify seeking epistemic or positional advantage vis-a-vis one’s interlocutor.

In fact, even without filling in such further details, we take it that there are significant advantages to be had by adopting this sort of Bergmanian thesis: that lying is wrong, when it is wrong, in virtue of the liar unjustly seeking either epistemic or a broader positional advantage. For one thing, it points to a way of explaining why political lies like Trump’s are wrong even when we all know that they are lies: because someone like Trump gains a positional advantage if he succeeds in increasing our credence that all politicians are liars. And, plausibly, it is unjustified to seek such an advantage in a modern democracy in general—particularly when your opponent actually is not (or is not much of) a liar.

References

- Audi, R. (1973). Intending. *Journal of Philosophy*, 70(13), 387–403.
- Augustine. (1952 [395]). Lying. In R. Deferrari (Ed.), *Treaties on various subjects* (Vol. 16, pp. 53–120). Washington, D.C.: Catholic University of America Press.
- Carson, T. (2006). The definition of lying. *Noûs*, 40, 284–306.
- Carson, T. (2010). *Lying and deception: Theory and practice*. Oxford and New York: Oxford University Press.

- Chisholm, R., & Feehan, T. (1977). The intent to deceive. *Journal of Philosophy*, 74, 143–159.
- Davidson, D. (1989 [1978]). Intending. In *Essays on actions and events* (pp. 83–102). Oxford: Clarendon Press.
- Fallis, D. (2009). What is lying? *Journal of Philosophy*, 106, 29–56.
- Fallis, D. (2015). Are bald-faced lies deceptive after all? *Ratio*, 28, 81–96.
- Faulkner, P. (2013). Lying and deceit. In H. LaFollette (Ed.), *International encyclopedia of ethics*. Hoboken, NJ: Wiley.
- Grice, H. (1973). Intention and uncertainty. *Proceedings of the British Academy*, 57, 263–279.
- Grice, H. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Keiser, J. (2015). Bald-faced lies: How to make a move in a language game without making a move in a conversation. *Philosophical Studies Online First*.
- Kissine, M. (2013). *From utterances to speech acts*. Cambridge and New York: Cambridge University Press.
- Krstić, V. (2019). Can You Lie Without Intending to Deceive? *Pacific Philosophical Quarterly*, 100, 642–660.
- Lackey, J. (2013). Lies and deception: An unhappy divorce. *Analysis*, 73(2), 236–248.
- Ludwig, K. (1992). Impossible doings. *Philosophical Studies* 65. 257–281.
- Mahon, J. (2011). Review of *Lying and Deception* by Thomas L. Carson. *Notre Dame Philosophical Reviews*.
- Maitra, I. (2018). Lying, Acting, and Asserting. In E. Michaelson and A. Stokke (eds.), *Lying: Language, Knowledge, Ethics, and Politics* (pp. 65–84). Oxford and New York: Oxford University Press.
- Meibauer, J. (2011). On lying: Intentionality, implicature, and imprecision. *Intercultural Pragmatics*, 8, 277–292.
- Meibauer, J. (2014). Bald-faced lies as acts of verbal aggression. *Journal of Verbal Aggression*, 2(1), 127–150.
- Mele, A. (2010). Intention. In T. O'Connor & C. Sandis (Eds.), *A companion to the philosophy of action* (pp. 108–113). Oxford: Wiley Blackwell.
- Neale, S. (2005). Pragmatics and binding. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 165–285). Oxford and New York: Oxford University Press.
- Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88, 251–264.
- Stokke, A. (2010). Intention-sensitive semantics. *Synthese*, 175, 383–404.
- Stokke, A. (2013). Lying and asserting. *Journal of Philosophy*, CX(1), 33–60.
- Stokke, A. (2018). *Lying and Insincerity*. Oxford and New York: Oxford University Press.

- Velleman, J. (1989). *Practical reflection*. Princeton, NJ: Princeton University Press.
- Williams, B. (2002). *Truth and truthfulness: An essay in genealogy*. Princeton, NJ: Princeton University Press.
- Williamson, T. (1996). Knowing and Asserting. *The Philosophical Review*, 105, 489-523.