Lies, Harm, and Practical Interests*

Andreas Stokke
Swedish Collegium for Advanced Studies
Uppsala University
andreas.stokke@filosofi.uu.se

Forthcoming in Philosophy and Phenomenological Research

Abstract

This paper outlines an account of the ethics of lying, which accommodates two main ideas about lying. The first of these, Anti-Deceptionalism, is the view that lying does not necessarily involve intentions to deceive. The second, Anti-Absolutism, is the view that lying is not always morally wrong. It is argued that lying is not wrong in itself, but rather the wrong in lying is explained by different factors in different cases. In some cases such factors may include deceptive intentions on the part of the liar. In other cases, where such intentions are not found, the wrong in lying may be explained by other factors. Moreover, it is argued that the interaction between considerations against lying and considerations against telling the truth are sensitive to the practical interests of those lied to. When the topic of the lie in question matters little to the victim's rational decision making, the threshold for when considerations against telling the truth can outweigh considerations against lying are lowered. This account is seen to explain why lying to avoid little harm is sometimes permissible, and sometimes not.

*Previous versions of this paper were presented at the Institute of Philosophy, London, the University of Bergen, and the University of Stockholm. I am grateful to the audiences on all of these occasions for useful discussion. I am particularly indebted to Don Fallis, Eliot Michaelson, Gunnar Björnsson, James Mahon, Jed Lewinsohn, Jesse Tomalty, Jessica Pepp, Kalle Grill, Ole Hjortland, and Pål Antonsen for very valuable suggestions.

1 Introduction

This paper outlines an account of the ethics of lying. The account is designed to accommodate two main ideas. The first is that, although lies are often told with the aim of deceiving those lied to, lying does not necessarily involve intentions to deceive. Call this *Anti-Deceptionalism*. The second is that, even though lying is often morally wrong, lying is not always morally wrong. Call this *Anti-Absolutism*.

Both Anti-Deceptionalism and Anti-Absolutism are untraditional claims about lying. On the one hand, a long tradition has characterized lying as saying what one believes to be false with the intent to deceive.[1] On the other hand, Absolutism – the view that lying is never morally justified – has notoriously been the position of a number of traditional views on the ethics of lying.[2] However, both Anti-Deceptionalism and Anti-Absolutism have been defended by more recent writers on lying.[3] These arguments will not be rehearsed here. Rather, the project for this paper is to ask how we can think of the ethics of lying *assuming* both Anti-Deceptionalism and Anti-Absolutism.

The view to be outlined here is one according to which lying is not wrong in itself, but rather the wrong in lying is explained by different factors in different cases. In some cases the relevant factor may be intentions to deceive. In other cases the wrong in lying derives from other circumstances. These may include harm, violation of rights, or other wrong-making factors. This will explain why non-deceptive lies are sometimes morally wrong.

Moral considerations against lying are often weighed against moral considerations against telling the truth. Perhaps most commonly, telling the truth may involve doing harm, and in such cases, considerations against doing harm may outweigh moral reasons not to lie. I will suggest that this kind of interaction between considerations against lying and considerations against telling the truth is sensitive to practical interests, understood, roughly, as what matters rationally for

[1] Proponents of this view include Augustine (1952 [395]b), (1952 [395]a), Isenberg (1964), Chisholm and Feehan (1977), Bok (1978), Kupfer (1982), Davidson (1985), Simpson (1992), Adler (1997), Williams (2002), Frankfurt (2005 [1986]), Faulkner (2007), (2013). An exception to this tradition is Aquinas (1922 [1265-74]) who does not include an intention to deceive in the definition of lying.

[2] Absolutists about lying include Augustine (1952 [395]a), Aquinas (1922 [1265-74]), Kant (1996 [1797]). For discussion, see, e.g., Paton (1954), Isenberg (1964), Bok (1978), Kupfer (1982), Korsgaard (1986), MacIntyre (1995), Williams (2002), Mahon (2006), (2009), Wood (2008), Shiffrin (2014).

[3] Proponents of Anti-Deceptionalism include Aquinas (1922 [1265-74]), Carson (2006), (2010), Sorensen (2007), Fallis (2009),Stokke (2013), Shiffrin (2014). Defenders of Anti-Absolutism include Mill (1979 [1863]), Sidgwick (1966 [1907]), Ross (2002 [1930]), Carson (2010), Shiffrin (2014).

decision making.

When the topic of the relevant lie has low bearing on the practical interests of those lied to, the threshold for when considerations against lying yield to considerations against telling the truth is lowered. The motivation behind this idea is that, when there is serious practical interest, knowing the truth about the relevant issue is important for the recipient's decision making, and this correlates with a higher pressure toward telling the truth, even if doing so may cause some harm. By contrast, when knowing the truth about the relevant issue is of low practical importance, it may be permissible to lie in order to avoid a lower degree of harm than it would otherwise be permissible to avoid by lying.

Section 2 argues for a view that accommodates Anti-Deceptionalism. I propose that lying is not wrong in itself, but is explained in terms of different factors, depending on the case. Section 3 extends the view to accommodate Anti-Absolutism by arguing that the interaction between considerations for and against lying, when they concern harm, is sensitive to practical interests.

2 Anti-Deceptionalism

2.1 Intrinsic, Non-Intrinsic, Unified, and Pluralist Views

Anti-Deceptionalism is motivated by examples of so-called *bald-faced lies*. That is undisguised lies that are not intended to deceive anyone. Consider, for instance, (a version of) Thomas Carson's (2006), (2010) much-discussed example of the Cheating Student.

> **The Cheating Student**
> A student accused of cheating on an exam is called to the Dean's office. The student knows that the Dean knows that she did in fact cheat. But it's also well known that the Dean won't punish someone who explicitly denies their guilt because, in such cases, taking action involves a great deal of bureaucracy. So the student says, "I didn't cheat."

Anti-Deceptionalists argue that while the student, in this case, is lying, she is not intending to deceive the Dean. As forewarned, the arguments for and against this claim are not under discussion here. I assume that examples like the Cheating Student establish that lies are sometimes told without intentions to deceive anyone.

Some critics of Anti-Deceptionalism argue that the view faces a challenge concerning the moral status of lying. Indeed, if all lies are intended to deceive, as the

traditional view holds, it is straightforward to explain the moral status of lying in terms of the moral status of deception. Roughly, lying will be wrong whenever deception is wrong. However, this type of account of the ethics of lying is not available to the Anti-Deceptionalist. Some find this a major cost of the view. For example, Jennifer Lackey (2013) writes,

> one natural criticism that we might have of the liar is that she is engaged in intentional deception, where such deceit carries the weight of the prima facie moral wrongness of such acts. Divorcing lying from deception, however, also divorces it from this explanation of its prima facie moral wrongness. (Lackey, 2013, 237)

Lackey takes this to be a central motivation for attempting to vindicate Deceptionalism about lying in the face of examples like the Cheating Student.

Nevertheless, the Anti-Deceptionalist is not without means of arguing for other ways of explaining the moral status of lying. First, the challenge Lackey presents is genuine only if some non-deceptive lies are morally wrong. Indeed, some Anti-Deceptionalists claim that non-deceptive lies, such as the student's lie to the Dean, are morally neutral. For example, Roy Sorensen (2007) has argued that while non-deceptive lies may be annoying, in bad taste, and may even be symptoms of immorality, they are morally neutral. If this is right, then there is no problem with Anti-Deceptionalism when it comes to explaining the moral status of lying. We could hold on to the claim that lying is wrong when deception is wrong.

However, many will find the claim that bald-faced lies are morally neutral hard to accept. I take it that we want to say that the student's lie to the Dean is morally wrong, and not just that it is improper for other reasons, even though it may be that as well. So in order to accommodate Anti-Deceptionalism, our account must accept that the wrong in lying cannot always be the wrong in deception, or the wrong in intended deception.

Yet it is important to note that, contrary to what is suggested by Lackey's objection, even the Deceptionalist may find reason to locate the wrong in lying somewhere else than in the deception she thinks invariably accompanies lying. There are two broad alternatives to the view that the wrong in lying is the wrong in deception. The first is the view that lying itself is a morally relevant factor, or as we might put it, lying itself is a moral wrong-maker. Call this view *Intrinsic*. Alternatively, one can hold that the wrong in lying derives from some factor or factors other than lying itself. Call this view *Non-Intrinsic*. To accommodate Anti-Deceptionalism, a Non-Intrinsic view must locate the wrong in lying, at least in some cases, elsewhere than in the moral status of deception.

We can note here that the choice between Intrinsic and Non-Intrinsic views is independent of the choice between Absolutism and Anti-Absolutism. It is of course perfectly possible for an Absolutist to adopt either an Intrinsic or a Non-Intrinsic view. One may think that, lying is wrong in itself and is never permissible, or indeed that even though the wrong in lying derives from some other factor or factors, it always so strong as to make lying impermissible.[4] On the other hand Anti-Absolutists may hold either Intrinsic or Non-Intrinsic views. One might think that even though lying is wrong in itself, it is sometimes permissible to lie, or that the wrong in lying derives from other factors, and sometimes these are not strong enough to make lying impermissible.[5]

It is reasonable to think that, from the point of view of theory-building, Non-Intrinsic views are preferable to Intrinsic views. The main reason for this is a general considerations in favor of parsimony. If the wrong in lying can be explained in terms of other factors we independently think are wrong, this is to be preferred over positing lying as a *sui generis* wrong-maker. Yet, it is important to note that any Intrinsic view is a potential way of accommodating Anti-Deceptionalism, and hence is a response to complaints such as Lackey's. If lying is wrong in itself, independently of its involvement of deception, then accepting that lying sometimes does not involve deception is in no way to give up an explanation of the wrong in lying.

In what follows we will focus on how to develop a plausible version of a Non-Intrinsic view of the wrong in lying. Such views come in two main varieties. A *Unified* Non-Intrinsic view holds that there is a single factor that explains the wrong in lying in every case. On such views, the ethics of lying may be derived from the ethics of the underlying factor. Indeed, a view such as the one Lackey laments the departure of, according to which the moral status of lies derives from the moral status of deception, is an example of a Unified Non-Intrinsic view. By contrast, a *Pluralist* Non-Intrinsic view holds that there are different factors that explain the wrong in lying in different cases.

In the rest of this section I will consider two candidates for a Unified Non-Intrinsic view, and I will suggest that they each require strong commitments in order to be plausible. For this reason, I will adopt the line that Pluralist Non-Intrinsic views are more promising. In the next section the task will be to argue for a way of squaring such a view with Anti-Absolutism.

[4]For example, at least on some readings, Kant's (1996 [1797]) notorious position is an instance of an Absolutist Intrinsic View.

[5]For example, an Anti-Absolutist Intrinsic view is held by W.D. Ross (2002 [1930]).

2.2 Manipulation and Coercion

One proposal for a Unified view stems from the widespread idea that the wrong in lying is the wrong in manipulation. To be consistent with Anti-Deceptionalism, this proposal requires accepting that manipulation does not necessarily involve deception, even if one thinks it often does. Don Fallis (2014) suggests a view of this kind:[6]

> [A] standard explanation of the *prima facie* wrongness of deceptive lies is that they are *manipulative* [...]. But many bald-faced lies are also manipulative. It is just that they do not achieve their manipulative effects by means of deception. (Fallis, 2014, 14)

One advantage of this proposal is that, as Fallis suggests, it is a weakening of a view of the wrong in deceptive lying. As such, it has the potential for explaining both cases. For example, Bernard Williams (2002) writes,

> In our own time we find it particularly natural to think deceiving people (or at least some people, in some circumstances) is an example of using or manipulating them, and that that is what is wrong with it. (Williams, 2002, 93)

And similarly, Sissela Bok (1978) suggests that the wrong in deceptive lying is the wrong in manipulation:

> Those who learn that they have been lied to in an important matter […] see that they were manipulated, that the deceit made them unable to make choices for themselves according to the most adequate information available, unable to act as they would have wanted, to act had they known all along. (Bok, 1978, 21-22)

However, contrary to Fallis's suggestion, the main problem for the view that the wrong in lying is the wrong in manipulation is that it is hard to square with Anti-Deceptionalism. A standard view is that manipulation is a species of deception, and that this is what distinguishes it from other ways of (attempting to) control or influence the will of another, such as coercion.[7] But if so, then given Anti-Deceptionalism, we cannot explain the wrong in lying as the wrong in manipulation. This conclusion is reached by Seana Shiffrin (2014):

[6] Fallis does not commit to a Unified view, but suggests that the wrongness of other kinds of bald-faced lies might be explained differently. I use his suggestion as an example of a possible Unified view.

[7] See, e.g., Todd (2013). But see also Baron (2003) for discussion.

> Many accounts of the wrong of the lie emphasize the wrong of manipulating or aiming to manipulate the will of the recipient; such accounts implicitly imagine that deception is an aim or product of the lie, for otherwise no such manipulation could occur. (Shiffrin, 2014, 22)

In other words, Shiffrin assumes that manipulation is a form of deception. Hence, since Shiffrin moreover accepts Anti-Deceptionalism, she concludes that the wrong in lying cannot be the wrong in manipulation.[8]

In light of this, another option is to try to explain the wrong in lying as *either* the wrong in manipulation or the wrong in coercion. After all, to the extent that coercion is understood, perhaps crudely, as non-deceptive manipulation, such a view would not be abandoning more than the letter of a Unified view.

However, some morally wrong lies are arguably not comfortably classified as coercive or manipulative at all. For example, it is natural to think that the student's lie to the Dean in the Cheating Student case is not an instance either of coercion or manipulation. According to one standard way of characterizing coercion,

> In cases of coercion, […] one tries to present one's target with considerations she will regard as irresistibly compelling compliance with one's demands. (Todd, 2013, 2)

Along the same lines, Allen Wood (2014) writes,

> I am coerced to do something when I either do not choose to do it or if, when I do choose to do it, I do it because I have no *acceptable alternative*. (Wood, 2014, 23)

But the Cheating Student case, as described, does not naturally fall into this category. The student's lie does not present the Dean with considerations that overwhelmingly favor not punishing her. Correspondingly, it is not right to say that the Dean has no acceptable alternative, at least when what constitutes an acceptable alternative is understood in the way that is arguably necessary for it to be a plausible characterization of coercion. Rather, the Dean could (and probably should) punish the student. It is just that the student has good reasons to think that she will not, due to the bureaucracy involved when a student has explicitly pleaded not guilty.

I does not look like a Unified view that appeals to manipulation or coercion is likely to be successful. There are arguably many lies that are morally wrong, but which are not plausibly described as involving either manipulation or coercion. Of course this does not rule out that there might be other ways of holding a Unified view of the wrong in lying. Next, I turn to an alternative proposal along these lines.

[8] See Shiffrin (2014, 13–14) for her rejection of Deceptionalism.

2.3 Harm

Another way of arguing for a Unified view is to propose that when lying is wrong it is because the lie is harmful to others. If this can be made plausible, it arguably has an immediate advantage, since it is independently plausible that we accept some kind of norm against harming others.

Shelly Kagan (1998) has pointed out that whether this style of account will work as a general account of the wrong in lying ultimately depends on our account of the nature of harm. Consider a version of a well-known type of example that Kagan adapts from Thomas Nagel (1979 [1970], 4):

> **The Deceived Businessman**
> Imagine a man who dies contended, thinking he has achieved everything he wanted in life: his wife and family love him, he is a respected member of the community, and he has founded a successful business. Or so he thinks. In reality, however, he has been completely deceived: his wife cheated on him, his daughter and son were only nice to him so that they would be able to borrow the car, the other members of the community only pretended to respect him for the sake of the charitable contributions he sometimes made, and his business partner has been embezzling funds from the company, which will soon go bankrupt. (Kagan, 1998, 34-35)

To make the case stronger, if necessary, one can imagine that the businessman's experiences, from the inside, are *exactly* as they would have been, had he not been deceived.

Since the lies told to the businessman are unquestionably morally wrong, to uphold the claim that the wrong in lying is invariably the wrong in harm, we would need to argue that the lies told to the businessman were harmful to him. One way of doing so is to accept a view according to which, roughly, you harm someone if you cause their well-being to be diminished, combined with some version of a preference-theory of well-being.[9] In particular, what is needed is a view according to which well-being is a matter, at least in part, of one's desires or preferences in fact being satisfied.[10] As Kagan notes, a suitable version of such a view would be in a position to count the lies told to the businessman as wrong by accepting that

[9] To be sure, we will need some view of how causing a decrease of well-being may be an instance of harm. But we will leave out discussion of this here.

[10] Most preference-theorists will endorse further constraints on when the satisfaction of a desire is valuable, e.g., depending on to what extent the preference is informed by knowledge, or the like. See, e.g., Brandt (1979). These details need not concern us here.

"lies can still harm someone even if they do not adversely affect his mental states." (Kagan, 1998, 109)

The first thing to note about this proposal is that not everyone will accept the theory of well-being that underlies it.[11] However, it is worth asking whether it might be successful, if we do accept such a view. Kagan suggests that "given a preference theory of well-being, it will be extraordinarily difficult – and perhaps impossible – to describe a case in which a lie is told but no harm is done." (Kagan, 1998, 110) The reason is, Kagan argues, that it is possible to point to preferences that arguably most people have and which are arguably thwarted by all lies:

> Most of us desire to be in honest and open relationships with others; this preference is directly thwarted when others lie to us, even if we realize they are lying. (Ibid.)

The last point is important, since it means that even non-deceptive lies might be harmful, independently of whether they are believed. Consider the student's lie to the Dean. To be sure, the Dean might not necessarily desire to be in an "honest and open" relationship with the student. The Dean might not really be interested in her relationship with the student in that sense. But still, it is plausible to think that the Dean prefers not being lied to. And hence, given the kind of view under consideration, the lie harms the Dean by thwarting this preference.

Yet some might feel that this does not quite get to the core of the matter. It does not seem that what is wrong with the student's lie is that it involves harm to *the Dean*. We might think of variations of the example where the student is required to testify by filling out a form, instead of being interviewed by the Dean. Such cases might be elaborated such that it is unclear whether anyone is the recipient of the lie. But nevertheless it is natural to think that, if nothing else, the student's lie thwarts the preferences of those who prefer that students do not lie about their exam practices, such as the administration, trustees, parents, or other students. And in particular, given the kind of preference-theory of well-being we are envisioning, such preferences may be thwarted even if those who have them do not find out about the lie. Just like the lies told to the unsuspecting businessman were said to harm him by thwarting his preferences, the student's lie might be said to thwart a general preference against lying, even among those who are unaware of the lie.

[11] For criticism of preference-theories of well-being, see, e.g., Parfit (1984, Appendix 1), Kraut (1994), Sobel (1994), Sumner (1996). For defenses, see, e.g., Brandt (1979), Murphy (1999), Bykvist (2002), Heathwood (2005).

However, one might begin to worry about this view from a different perspective at this point. First, to vindicate a Unified Non-Intrinsic view of the wrong in lying, we adopted a controversial view of the factor – i.e., harm – that is meant to play the unifying role. But moreover, to uphold the claim that all morally wrong lies thwart some preference or other, we have been forced to posit a widespread and general preference against lying itself. That is, we argued that lying may be harmful because it can thwart a preference against lying itself, even when no one is aware of it.

Even though it is true that this is not a version of an Intrinsic view, that is, the wrong in lying is not explained by supposing that lying is wrong in itself, it might start to look like what was meant to be the attraction of Non-Intrinsic views is receding. In particular, the advertised advantage was parsimony, and more specifically, the appeal of explaining the wrong in lying in terms of factors we independently think are morally relevant, in this case, harm. But if the view is ultimately forced to accept that harm may arise simply because people prefer not being lied to, we have not achieved much in terms of this kind of theoretical desideratum.

Considered alongside the fact that the view we have sketched requires a commitment to controversial theories of harm and well-being, we might think it is more promising to opt for a Pluralist version of a Non-Intrinsic view of the wrong in lying.

2.4 Pluralism

According to Pluralist Non-Intrinsic views, for each case in which lying is wrong, the wrong is explained in terms of some factor other than the lying itself, although not necessarily the same factor in every case. The first thing to note is that if we adopt a Pluralist view, we can accept the idea behind the kind of complaint we saw Lackey giving voice to earlier. Namely, the relevant factor in why it is wrong to lie might be deception in a number of cases. For some Anti-Deceptionalists, this will already be an attraction of Pluralist over Unified views. If we are Pluralists, we can easily accept that the wrong in lying is often the wrong in deception – perhaps even in all cases of deceptive lying. Yet we will be in a position to accept that, in some cases, the wrong in lying is something else. For example, we might think that what is wrong with the student's lie to the Dean is that it obstructs justice, or something else.

The Pluralist's view is compatible with there being a limited range of factors that are responsible for the wrong in lying in all (or most) cases. For instance, a Pluralist might hold that at least a great many cases can be explained in terms of

either the wrong in harm or in terms of some violation of rights. In particular, some might argue that the wrong in obstructing justice, as well as the wrong in manipulation and coercion, ultimately derive from a violation of rights.[12]

We will forego the daunting task of surveying – let alone assessing – the range of possible views here. Instead, we will proceed under the assumption that the most plausible candidate is an account that explains the wrong in lying in terms of a, possibly limited, range of different factors for different cases. As we have seen such a view can accommodate Anti-Deceptionalism. In the next section we consider how it fares with respect to Anti-Absolutism.

3 Anti-Absolutism

3.1 Weighing for and against Lying

Given a Pluralist view of the kind we have outlined, in most situations some considerations will morally favor not telling a particular lie. For example, in many cases considerations concerning the harm that will result from telling a lie may favor not telling the lie. In other cases considerations concerning some violation of rights that will be involved in telling a lie may favor not telling the lie. And in still other cases there may be other considerations or some combination of considerations.

Correspondingly, one approach to Anti-Absolutism is to think that, in cases where lying is not wrong, considerations against lying are outweighed by other norms or considerations, whereas in cases where lying is wrong, considerations against lying outweigh other norms or considerations, if any. To be sure, it is at least logically possible that there are also cases in which lying is not wrong because there simply are no considerations against it. However, we are interested here in cases where there are at least some initial moral considerations against lying.

What kind of other considerations or norms can outweigh considerations against lying? Presumably, the list is at least as long as the list of factors that might go into considerations against lying in the first place. Yet, since we have not decided on the number of factors, we cannot consider the full range of cases, even in the abstract. Moreover, doing so would be infeasible. Instead, we will narrow our perspective to considering just one factor, harm. We confine ourselves, therefore, to cases in which considerations against lying are weighed in relation to considerations concerning some harm that is expected to result from telling the truth. Such cases are

[12]On justice in this regard, see, e.g., Kagan (1998, 176–177) for discussion. For some related discussion of manipulation and coercion, see Wood (2014).

arguably very common; indeed, they are arguably the most common kind of case that have been discussed in the literature on Anti-Absolutism.

So the question before us is, under what circumstances can considerations against harm outweigh considerations against lying? When thinking about this question it is important not to prejudge other matters. In particular, we will not assume any particular theory of harm or of doing harm, such as the one involving preference-theories of well-being discussed earlier. Rather, we will take a broader, non-committal view of harm in examining the ways in which considerations against harm may interact with considerations against lying.

We will look at two categories of examples from the canon. The first are cases in which, according to the Anti-Absolutist, lying is not morally wrong. We will consider the well-known examples of the Murderer at the Door and the Dying Old Woman.

> **Murderer at the Door**
> A murderer comes to your door and asks for the whereabouts of her victim. You know the victim is in her office. You tell the murderer, "She's at the mall."

> **Dying Old Woman**
> An old woman on her deathbed asks you, "How is my son?" You know the son was badly hurt in a car-accident the day before. You tell her, "He's happy and healthy."

The second category are cases in which lying is wrong. We will consider the following two examples:

> **Affair**
> Your spouse confronts you with the question of whether you are having an affair. You are. But you answer, "No, I'm not having an affair."

> **The Cheating Student**
> A student accused of cheating on an exam is called to the Dean's office. The student knows that the Dean knows that she did in fact cheat. But it's also well known that the Dean won't punish someone who explicitly denies their guilt because, in such cases, taking action involves a great deal of bureaucracy. So the student says, "I didn't cheat."

The strategy, then, will be to consider why it is wrong to lie in the Affair case and the Cheating Student case, while it is not wrong to lie in the Dying Old Woman case

and the Murderer at the Door case. And in particular, we are after an explanation in terms of the considerations against lying being outweighed by considerations concerning harm that is expected to result from telling the truth.

3.2 The Murderer and the Affair

Why is it wrong to lie to one's spouse about an affair while it is not wrong to lie to the murderer at the door? One suggestion is as follows. In both cases the lie will prevent harm that would result from telling the truth. In the Affair case lying will prevent harm your spouse will suffer as a result of being told about the affair. She will be saddened, insulted, and so on. In the Murderer case lying will prevent a loss of life. Yet, so the suggestion goes, in the Affair case, the harm that can be prevented by lying is not great enough to outweigh the relevant considerations against lying.

To be sure, we are not ruling out that there might be cases in which one has reason to think that telling one's spouse the truth about having had an affair is likely to cause a very high degree of harm, and hence in such cases, it might be thought morally permissible to lie in order to prevent such an outcome. For example, if you have reason to think that your spouse will kill herself, or your lover, if you tell her about the affair, then most likely, it will be seen as morally permissible to lie. Still, it is plausible to think that there are cases in which the degree of harm that will result from telling the truth is insufficient to outweigh considerations against lying.

By contrast, in the Murderer case, it is natural to think that the harm that can be prevented by lying is great enough to outweigh considerations against lying. We allow ourselves to be neutral on what the considerations against lying amount to in these cases. All that is assumed is that, as seems undeniable, some moral considerations or other at least initially speak against lying in each of these cases.

We should note here that, on some views (e.g., Korsgaard, 1986, Shiffrin, 2014), the fact that the recipient of the lie is a would-be murderer is a morally relevant factor in itself. Briefly, on such views, the fact that you know (or have good reason to believe) that the recipient is likely to use your information in order to act in an evil way means that you have no obligation at all to tell her the truth. By contrast, on the view we are considering here, the relevant feature of the case is that, by lying, you can prevent a degree of harm that is sufficient to permit lying.

Comparing cases like the Murderer at the Door and the Affair provides a *prima facie* reason to think that the threshold for when considerations against lying are outweighed by considerations against harm is by default relatively high. That is,

typically considerations against lying are outweighed by considerations concerning harm that might result from telling the truth only when the harm that can be prevented by lying is quite significant. In other words, a simple explanation of the contrast between the Murderer and the Affair cases is to point to such a high threshold for when lying is permitted by the possibility of preventing harm that would result from telling the truth.

An alternative proposal is to point to the relation between the parties in the two cases. In particular, it might be suggested that the special relation you have to your spouse gives extra weight to considerations against lying. It is not unreasonable to think that the reason it is impermissible to lie to one's spouse about an affair is not just that the degree of harm that might be prevented by doing so is not great enough. Rather, it might be thought that at least part of the reason has to do with a special obligation one has to be truthful to an intimate relation.

To test this suggestion, imagine that you know that the wife of a not close acquaintance – someone toward which you do not have any relevant special obligations – is having an affair. Suppose this acquaintance asks you whether you think his wife is having an affair. Most people would think it is wrong to lie in reply. If so, then this suggests that, in the original Affair case, the potential harm in telling the truth is not sufficient to outweigh considerations against lying, independently of the special relation between you and your spouse. In particular, we can arguably construct cases in which the degree of harm that would be involved in telling the truth in both cases is the same.

So we have reason to think that the wrong in lying in the Affair case is better explained as the potential harm in telling the truth being insufficient to outweigh considerations against lying than in terms of the weight of considerations against lying being increased by a special relation. In turn, therefore, this reinforces the suggestion that the threshold for when considerations against harm outweigh considerations against lying is by default relatively demanding.

The challenge for the Anti-Absolutist arises due to the fact that there are cases in which it seems to be permissible to lie in order to prevent a lower degree of harm than would otherwise be sufficient to outweigh considerations against lying. This suggests that the threshold for when considerations against harm can outweigh considerations against lying is sensitive to features of the situation. Below I sketch a way of understanding this kind of flexibility.

3.3 Practical Interests

Consider the question of why it is not wrong to lie to the Dying Old Woman. Given the above, a perhaps natural idea is that, in this case, the degree of harm that would result from telling her the truth about her son is sufficient to outweigh considerations against lying. After all, since the old woman presumably cares a great deal about her son's well-being, it is not unreasonable to think that the reason we consider it permissible to lie, in this case, stems from the arguably intense degree of pain the truth would cause her.

However, there are reasons to think that this cannot be right as it stands. Suppose the question was instead whether her husband had been faithful to her. Most will think that lying is still permissible. But in this scenario, or in some suitable version of it, arguably the same degree of potential harm is involved as in the Affair case – or, depending on your relation to the dying old woman, as in the case of the wife of your acquaintance having an affair. Similarly, we do not think that it is permissible to lie to old women who are not dying about the well-being of their sons, even when we realize that telling the truth will cause sadness and suffering.

This suggests that some special circumstance in the Dying Old Woman case is responsible for the fact that considerations against lying are outweighed by considerations against harm, even though the degree of harm that might be prevented by telling the truth would otherwise not be sufficient to permit lying. Indeed, this is what strikes us about the case at first sight. The proposal I want to explore here is that what is special about these kinds of cases has to do with practical interests.

To illustrate the motivation for this idea, it is useful to consider the familiar case of polite lies. Here is an example:

> **Newly Coiffed**
> Your friend asks you, "Do you like my new hairdo?" You don't like it. But you answer, "Yeah, it looks good."

Why is permissible to lie in this kind of situation? I suggest that the reason is that the threshold for when considerations against lying yield to considerations against harm is lowered in cases of this kind. That is, some feature of the situation is responsible for the fact that considerations concerning the harm that would result from telling the truth outweigh the active considerations against lying, whatever they are, even though they would otherwise not have been outweighed.

What is this feature? Cases involving polite lies suggest that a central factor in how such thresholds are fixed is the *topic* of the lie in question. That is, whether a lie is wrong or not depends, among other things, on what the lie is *about*. In

particular, it is natural to think that a lot depends on whether what the lie is about *matters* to the recipient. But matters in what way? The suggestion here will be that the strength of considerations against lying *vis-à-vis* considerations against doing harm by telling the truth is sensitive to the bearing of the topic of the lie on the practical interests of the recipient of the lie.

To spell out this proposal we need some idea of practical interests. A standard view is that what is of practical interest to an agent is what she ought rationally to consider in decision making. For example, Jason Stanley (2005) characterizes the notion as follows:

> A subject's interests determine her goals. Given these goals, there will be a range of actions which that subject ought practically to consider. Given that we are not ideal rational agents, there will be a range of alternatives that it will be legitimate to ignore. The rest of the alternatives to her beliefs are ones that she ought rationally to consider. A proposition is a serious practical question for an agent, if there are alternatives to that proposition that the subject ought rationally to consider in decision making. (Stanley, 2005, 92)

So a proposition is said to be of serious practical interest when one ought to consider its alternatives in decision making.[13] The alternatives to a proposition p can be seen as the range of candidate answers to some salient question to which p is an answer. For simplicity, we can confine ourselves here to polar, i.e., "yes-no," questions. Ultimately a theory of this kind should also consider wh-questions. In many cases, if it is of practical interest to one whether Mary is working, one ought rationally to consider not just the possibilities that Mary is working and that she is not, but also a range of other propositions like *Jim is working*, *Sarah is working*, *Only Mary is working*, *Jim and Mary are working*, and so on.[14] However, we will focus here on polar questions like *Is Mary working?* and the corresponding set of alternatives, i.e., *Mary is working* and *Mary is not working*. So, to take another example, the question of whether there is a 3.30 train might be of serious practical interest to me because, in deliberating about how to get to town in time for a meeting, I ought rationally to consider both the possibility that there is a 3.30 train and the possibility that there is not.

Given this, we can formulate a proposal as follows.

Threshold Principle
Let p be the content of the lie under consideration. Then

[13] See also Joyce (2002) for discussion.

[14] I defend a view of lying and what is said in terms of a theory of questions in Stokke (2016).

> If $p?$ is not of serious practical interest to the recipient of the lie, the harm-threshold for considerations against lying by asserting p is lowered.

Here "$p?$" denotes the polar question whether p, and the harm-threshold for considerations against lying is the threshold for when considerations against harm outweigh considerations against lying. So the principle says that if the question of whether the content of a particular lie is true is not something that the recipient ought rationally to consider in decision making, telling that lie is permitted in order to prevent a lower degree of harm than would otherwise be needed to outweigh considerations against lying.

This principle reflects the observation that, in many situations, one is weighing for and against lying about an issue for which knowing the truth of the matter is of practical importance for the recipient. In such situations it is natural to think that there is more pressure toward telling the truth, even if it will cause harm. Conversely, when knowing the truth of the matter is of less practical significance, considerations against telling the truth because of the potential harm involved in doing so are given more weight. When some issue matters for the decision-making of our interlocutors, we are more stringently required not to lie about it, even if it causes harm, whereas when the topic of conversation is of lower practical interest, we are more liberally permitted to lie in order to avoid harm.

It is worth stressing that this proposal is intended only to apply to interactions between considerations against lying and considerations concerning harm. In particular, I am not suggesting that the weighing of other relevant factors is also sensitive to practical interests. Indeed, it is plausible to think that there are factors for which practical interests do not matter. For example, suppose you think that people have the right to know who their parents are. If so, you might also think that it is impermissible to lie to someone about who their parents are, even if the question of who their parents are is not of serious practical interest to them. The suggestion here, then, is merely that in the – arguably, wide – range of cases when considerations against lying are weighed in relation to considerations against harm, the threshold for the latter outweighing the former is affected by practical interests.

3.4 Outweighing Considerations against Lying

We can now explain the permissibility of the polite lie in the Newly Coiffed example. Since we may assume that the question of whether you like your friend's hairdo is not of serious practical interest to your friend, it is permissible to lie in

order to avoid the harm that would result from telling her the truth. Note that, of course, this is perfectly consistent with thinking that the question of your opinion about the hairdo matters to your friend in other ways. Most likely, the fact that it matters in other ways is precisely the reason you think she will be hurt by the truth. The claim is just that whether you like the hairdo is not a question that your friend, in this case, ought rationally to consider in decision making.

In other words, according to the view we are sketching, the reason it is permissible for you to lie about the hairdo is that your friend does not have any goals such that reaching them rationally requires her to practically consider the question of whether you like her new hairdo or not. Of course, there may be other situations in which this is not so. For example, suppose your friend is going to have her wedding pictures taken and she wants her hair to look good on the pictures. If the question comes up in a situation of this kind, your opinion about the hair may matter for your friend's decision making. If so, the harm-threshold for considerations against lying remains high, and it would be wrong to lie in order to avoid the harm that would result from telling her you do not like the hairdo. I take it that this is the result we want.

To be sure, ultimately, we will need to refine the Threshold Principle in various ways. For one thing, it is likely that it will need to be relativized to the beliefs of the recipient. For example, imagine someone who irrationally believes that some question is of serious practical interest. A religious fundamentalist might irrationally believe that whether she should accept a blood transfusion depends strictly on whether the donor belongs to the same religious group as herself. We might think that it would be wrong to lie to her about whether the blood came from such a donor. Indeed, some might think that lying would be wrong even though a great deal of harm could be prevented, for instance, severe illness or perhaps even the loss of the patient's life. But since the question is not *in fact* of practical interest, the principle above permits lying about the donor. In light of this, we might want to adopt a modified version of the principle along the following lines:

> **Threshold Principle** (Relativized)
> Let p be the content of the lie under consideration. Then
> If S believes that p? is not of serious practical interest, the harm-threshold for considerations against lying to S by asserting p is lowered.

However, we will ignore this extra complication and consider the unrelativized principle.

Equipped with this kind of view of the interaction between considerations against lying and considerations against harm, we can explain the Dying Old Woman ex-

ample. Whether her son is alive is not of serious practical interest to the dying woman. She has no decisions to make for which she ought rationally to consider whether her son is alive or not. Therefore, the threshold is lowered, and it is permissible to avoid harm by lying.

As before, this is consistent with thinking that the question matters greatly in other respects. Indeed, that is why you think she will be saddened if you tell the truth. And moreover, this account is likewise consistent with there being otherwise similar situations in which the question of her son's well-being is of serious practical interest to the dying old woman. In such cases, the account does not permit lying to the dying old woman, which seems to be the right result.

We also explain the Affair case. The question of whether you are having an affair is of serious practical interest to your spouse. Hence, the harm-threshold is not lowered, and it is therefore impermissible to lie in order to avoid the degree of harm that will result from telling the truth. The same applies to the case of the wife of your not-close acquaintance having an affair. Since the harm-threshold is not lowered, due to serious practical interest, lying is not permitted.

Now consider the Murderer at the Door example. The question of the whereabouts of the would-be victim is of serious practical interest to the murderer. Hence, according to our account, the harm-threshold is not lowered. But nevertheless, of course the degree of harm that can be prevented by telling the lie, in this case, is high enough that we do not need the threshold to be lowered in order for considerations against harm to outweigh considerations against lying. In other words, the degree of harm that can be prevented by lying is high enough to permit lying, even though there is serious practical interest.

Finally, consider the Cheating Student example. It is plausible to think that if any harm can be avoided by telling the lie, it is harm to the student herself. Hence, it is not clear that moral considerations against harm are actively weighing against telling the truth at all in this case. That is, we might think that the considerations against harm that are relevant in cases of weighing moral reasons for and against lying have to do with harm against (innocent) others. But even if there are active moral considerations against harm, since the question of whether the student cheated is of serious practical interest, the threshold is not lowered, and arguably the degree of harm that can be avoided by lying is not great enough to permit lying.

4 Conclusion

Although lying is often morally wrong, lying is sometimes morally permissible. The most plausible shape for an account of this moral status of lying is a Non-Intrinsic Pluralist view, according to which lying is not wrong in itself, but rather the wrong in lying is explained by different factors in different cases. Since these factors may or may not include deceptive intentions on the part of the liar, this kind of view accommodates the Anti-Deceptionalist observation that lying does not always involve intentions to deceive.

Sometimes it is permissible to lie to avoid a degree of harm that in other cases appears to be to low to outweigh considerations against lying. The reason is that our weighing of considerations for and against lying is sensitive to the practical interest of our addressees. If the topic of the lie in question is something that matters little for the addressee's rational decision making (or if they believe that it does), the threshold for when considerations against lying can be outweighed by, for example, considerations against harm is lowered.

References

Adler, J. (1997). Lying, deceiving, or falsely implicating. *Journal of Philosophy*, *94*, 435–452.

Aquinas, T. (1922 [1265-74]). Of lying. In *Summa theologica* (Vol. XII, pp. 85–98). London: Burns, Oates, and Washbourne.

Augustine. (1952 [395]a). Against lying. In R. Deferrari (Ed.), *Treaties on various subjects* (Vol. 16, pp. 125–179). Washington, D.C.: Catholic University of America Press.

Augustine. (1952 [395]b). Lying. In R. Deferrari (Ed.), *Treaties on various subjects* (Vol. 16, pp. 53–120). Washington, D.C.: Catholic University of America Press.

Baron, M. (2003). Manipulativeness. *Proceedings and Addresses of the American Philosophical Association*, *77*, 37–54.

Bok, S. (1978). *Lying: Moral choice in private and public life*. New York: Random House.

Brandt, R. (1979). *A theory of the right and the good*. Oxford and New York: Oxford University Press.

Bykvist, K. (2002). Sumner on desires and well-being. *Canadian Journal of Philosophy*, *32*(4), 475–490.

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. (1985). Deception and division. In J. Elster (Ed.), *The multiple self* (pp. 79–92). Cambridge and New York: Cambridge University Press.

Fallis, D. (2009). What is lying? *Journal of Philosophy*, *106*, 29–56.

Fallis, D. (2014). Are bald-faced lies deceptive after all? *Ratio Online*, 1–16.

Faulkner, P. (2007). What is wrong with lying? *Philosophy and Phenomenological Research*, *75*, 535–557.

Faulkner, P. (2013). Lying and deceit. In H. LaFollette (Ed.), *International encyclopedia of ethics*. Hoboken, NJ: Wiley.

Frankfurt, H. (2005 [1986]). *On bullshit*. Princeton, NJ: Princeton University Press.

Heathwood, C. (2005). The problem of defective desires. *Australasian Journal of Philosophy*, *83*, 487–504.

Isenberg, A. (1964). Deontology and the ethics of lying. *Philosophy and Phenomenological Research*, *24*, 465–480.

Joyce, J. (2002). Levi on causal decision theory and the possibility of predicting one's own actions. *Philosophical Studies*, *110*, 69–102.

Kagan, S. (1998). *Normative ethics*. Boulder, CO: Westview Press.

Kant, I. (1996 [1797]). On a supposed right to lie from philanthropy. In M. Gregor (Ed.), *Immanuel Kant, practical philosophy* (pp. 611–615). Cambridge and New York: Cambridge University Press.

Korsgaard, C. (1986). The right to lie: Kant on dealing with evil. *Philosophy and Public Affairs*, *15*, 325–349.

Kraut, R. (1994). Desire and the human good. *Proceedings and Addresses of the American Philosophical Association*, *68*, 39–54.

Kupfer, J. (1982). The moral presumption against lying. *Review of Metaphysics*, *36*, 103–126.

Lackey, J. (2013). Lies and deception: An unhappy divorce. *Analysis*, *73*(2), 236–248.

MacIntyre, A. (1995). Truthfulness, lies, and moral philosophers: What can we learn from Mill and Kant? *The Tanner Lectures on Human Values*, *16*, 307–

361.
Mahon, J. (2006). Kant and the perfect duty to others not to lie. *British Journal for the History of Philosophy*, *14*, 653–685.
Mahon, J. (2009). The truth about Kant on lies. In C. Martin (Ed.), *The philosophy of deception* (pp. 201–224). Oxford and New York: Oxford University Press.
Mill, J. (1979 [1863]). *Utilitarianism*. Indianapolis: Hackett.
Murphy, M. (1999). The simple desire-fulfillment theory. *Noûs*, *33*, 247–272.
Nagel, T. (1979 [1970]). Death. In *Mortal questions* (pp. 1–10). Cambridge and New York: Cambridge University Press.
Parfit, D. (1984). *Reasons and persons*. Oxford and New York: Oxford University Press.
Paton, H. (1954). An allegeded right to lie: A problem in Kantian ethics. *Kant-Studien*, *45*, 190–203.
Ross, W. (2002 [1930]). *The right and the good* (P. Stratton-Lake, Ed.). Oxford and New York: Oxford University Press.
Shiffrin, S. (2014). *Speech matters: On lying, morality, and the law*. Princeton, NJ: Princeton University Press.
Sidgwick, H. (1966 [1907]). *The methods of ethics*. New York: Dover.
Simpson, D. (1992). Lying, liars and language. *Philosophy and Phenomenological Research*, *52*(3), 623–639.
Sobel, D. (1994). Full information accounts of well-being. *Ethics*, *104*, 784–810.
Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, *88*, 251–264.
Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
Stokke, A. (2013). Lying and asserting. *Journal of Philosophy*, *CX*(1), 33–60.
Stokke, A. (2016). Lying and misleading in discourse. *Philosophical Review*, *125*(1), 83–134.
Sumner, W. (1996). *Wellfare, happiness, and ethics*. Oxford: Clarendon.
Todd, P. (2013). Manipulation. In H. LaFollette (Ed.), *The international encyclopedia of ethics* (pp. 3139–3145). Oxford: Blackwell.
Williams, B. (2002). *Truth and truthfulness: An essay in genealogy*. Princeton, NJ: Princeton University Press.
Wood, A. (2008). *Kantian ethics*. Cambridge and New York: Cambridge University Press.
Wood, A. (2014). Coercion, manipulation, and exploitation. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and practice* (pp. 17–51). Oxford and New York: Oxford University Press.