

# The Truth About Deception<sup>1</sup>

JAPA PALLIKKATHAYIL

*University of Pittsburgh*

The prohibition on lying is often thought to be very stringent. Some have even been tempted to think that it is absolute. In contrast, the prohibition on other forms of deception seems to be looser. This paper explores the relationship between the duty not to deceive and the duty not to lie. This discussion is situated in the context of a broadly Kantian account of morality. Kant himself infamously claimed that one ought not lie to a murderer at the door about the location of his intended victim. This paper aims to explain how a broadly Kantian view can endorse a distinctive duty not to lie without thereby being committed to this kind of conclusion.

The prohibition on lying is often thought to be very stringent. Some have even been tempted to think that it is absolute. In contrast, the prohibition on other forms of deception seems to be looser. This paper explores the relationship between the duty not to deceive and the duty not to lie. Many people have investigated the apparent differences between these duties by trying to explain, or explain away, the special status of lying.<sup>2</sup> In this paper, approach the relationship between these duties from the opposite direction. I begin by giving an account of the duty not to deceive. And then with this account in place, I consider the how the duty not to lie might differ.

This discussion is situated in the context of a broadly Kantian account of morality. Kant himself infamously claimed that one ought not lie to a murderer at the door about the location of his intended victim.<sup>3</sup> This paper aims to explain how a broadly Kantian

---

<sup>1</sup> I received helpful feedback on earlier versions of this article from audiences at Union College; the “What is Wrong With Deception?” conference at the University of Reading; the Workshop on Authority, Coercion and Paternalism at the University of Graz; York College; SUNY Albany; the Fordham Epistemology and Ethics Workshop; Simon Fraser; the Midsummer Philosophy Workshop; the Problems in Legal Philosophy seminar at Columbia University; and the University of Pittsburgh. I am also especially indebted to Julia Driver, Grace Helton, Hallie Liberto, Joseph Raz, and Nishi Shah for helpful comments.

<sup>2</sup> For a very helpful overview of the attempts to specify the moral objection to lying, see chapter four of Jennifer Saul, *Lying, misleading, and what is said: an exploration in philosophy of language and in ethics* (Oxford: Oxford University Press, 2012).

<sup>3</sup> A great deal has been written about Kant’s claim about lying to the murderer at the door. For a helpful discussion of how this claim should be interpreted, see Allen W. Wood, “Lies,” in *Kantian Ethics* (New York: Cambridge University Press, 2008). For a prominent Kantian attempt to avoid Kant’s conclusion, see Christine M. Korsgaard, “The right to lie: Kant on dealing with evil,” in *Creating the Kingdom of Ends* (Cambridge, Cambridge University Press, 1996). My project in this paper differs significantly from Korsgaard’s in that I make no attempt to connect my discussion directly with Kant’s views about the Categorical Imperative. My discussion makes use of Kantian themes that might eventually be connected more explicitly with this element of Kant’s views. But I make no attempt to do so here.

view can endorse a distinctive duty not to lie without thereby being committed to this kind of conclusion.

My investigation of the relationship between the duty not to deceive and the duty not to lie is complicated by a debate about the definitions of deception and lying. It is tempting to think that the moral questions surrounding deception and lying cannot be answered prior to settling the nature of these activities.<sup>4</sup> In the course of giving my account, I will show that this temptation is a mistake. My account will provide a way of navigating borderline cases of deception and lying without settling definitions of these categories.

## 1. The Duty Not to Deceive

I am going to argue that the relationship between the duty not to deceive and the duty not to lie is analogous in an important way to the relationship between the duty of beneficence and the duty to keep one's promises. The key element in this analogy is the abstract structure of the duties involved. The duty of beneficence is an imperfect duty, i.e. a duty that is, in the first instance, a duty to adopt a particular end rather than to perform particular actions. In contrast, the duty to keep one's promises is a perfect duty, i.e. a duty to perform or refrain from particular actions. In this section, I develop the analogy between the duty not to deceive and the duty of beneficence. In Section 2, I develop the analogy between the duty not to lie and the duty to keep one's promises.

### 1.1. The Duty of Beneficence

In order to prepare the way for the analogy between the duty not to deceive and the duty of beneficence, I will begin with a brief description of the duty of beneficence. While Kant sometimes frames the duty of beneficence with reference to the happiness of others, this duty is best understood as focused on taking another's ends as reason giving, even if these ends come at the cost of her happiness. A person may, for example, self-sacrificially choose to take a thankless job. The appropriate reaction in this case is not to act to thwart this person's choice in order to secure her happiness or well being, but rather to take her choice as reason giving.<sup>5</sup> This is not to deny that a person's happiness is a source of reasons for us. One might, for example, have reason to give another a present that does not further any of her particular ends but makes her happy or contributes to her well being. I take it, though, that at least in our interactions with fellow competent adults, our reasons to promote their happiness are constrained by their ends in the way suggested by the thankless job example. In keeping with this understanding of the duty of beneficence, I am going to talk about taking others' ends as reason giving and set aside talk of happiness.

---

<sup>4</sup> Consider: "Conceptual questions about the nature of lying are prior to questions about the moral status of lying. Any theory about the moral status of lying presupposes an account of what lying is." Thomas L. Carson, "The Definition of Lying," *Noûs* 40:2 (2006), 284. Don Fallis also approvingly cites Carson's claim. See Don Fallis, "What Is Lying," *The Journal of Philosophy* 106:1 (2009), 30. For a very helpful overview of the debate regarding the definitions of deception and lying, see James Edwin Mahon, "The Definition of Lying and Deception", *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/lying-definition/>>.

<sup>5</sup> This example comes from Kyla Ebels-Duggan, who offers a compelling argument for this diagnosis of our duties in this case. See Kyla Ebels-Duggan, "Against Beneficence: A Normative Account of Love," *Ethics*, Vol. 119, No. 1 (October 2008) 152.

Note three important features of the duty of beneficence so understood. First, as I have indicated, the duty of beneficence is an imperfect duty. While perfect duties are duties to perform or refrain from particular actions, imperfect duties are, in the first instance, duties to adopt a particular end. So, for example, Kant argues that there is a perfect duty not to commit suicide and an imperfect duty to develop one's own talents. Adopting the end of developing one's talents does not typically require one to engage in any specific action. This end alone, for example, does not tell in favor of piano rather than violin lessons. So, in general, we have discretion about how to satisfy an imperfect duty.

Nonetheless, there may be cases in which an imperfect duty decisively requires one course of action. As Thomas Hill points out, "in dire straits an unfortunate person might have only one (permissible) chance to escape a debilitating, brain-numbing life of physical labor. Given background conditions, taking the chance might be strictly required because nothing else would count as having seriously made developing her talents an end."<sup>6</sup> Likewise, although we often have some discretion about how to satisfy the imperfect duty of beneficence, there are cases in which one must act in a specific way in order to count as valuing another agent's ends at all, e.g. throwing a life preserver to a drowning man.<sup>7</sup>

The second noteworthy feature of the duty of beneficence is that it involves reasons both to help and not to interfere.<sup>8</sup> Suppose you encounter a couple taking photos of themselves in front of a famous monument. If you take their end as reason giving, you will take yourself to have reason not to interfere by, say, walking through their shot. You will also take yourself to have reason to help them by, say, taking a photo of them together. Third, notice that whether these reasons give rise to obligations not to interfere or to help depends to some extent on how acting in these ways may affect the other ends in play, including your own. If a person is rushing to make an important appointment, she is not required to stop and help the couple and may even be justified in running through their shot. In intimate relationships, decisions about how to respond to the ends of each participant may need to be made jointly. But the person rushing to make her appointment does not need to consult with the couple to weigh her ends against theirs; she is permitted to make the assessment herself, though there may be bounds on what could count as a legitimate assessment. Suppose a person is wandering about doing nothing in particular, but feels that staying out of the couple's shot is just not worth his effort. It seems that this person would be making a mistake akin to failing to throw the drowning man a life preserver.<sup>9</sup>

---

<sup>6</sup> Thomas Hill, "Meeting Needs and Doing Favors," *Human Welfare and Moral Worth: Kantian Perspectives* (Oxford: Oxford University Press: 2002) 208.

<sup>7</sup> For a helpful further discussion of this matter, see Karen Stohr, "Kantian Beneficence and the Problem of Obligatory Aid," *Journal of Moral Philosophy* 8 (2011) 45–67.

<sup>8</sup> The reasons to help and not to interfere are best described as merely *prima facie* reasons because they might be defeated by existing duties and not simply outweighed by competing ends. So, for example, my promise to help someone else might defeat and not merely outweigh the reason I would otherwise have had to help you. Since, however, this complication does not play a role in the argument, I leave it aside here.

<sup>9</sup> My claim here is that both failing to help the drowning man and not attempting to avoid walking through the couple's photo involve ignoring the reasons that stem from others' ends. It is compatible with this claim that failing to help the drowning man is, in some sense, a worse mistake because it involves ignoring more serious ends.

Should one treat *all* of the ends that others have as reason giving? What about immoral ends or ends that seem for all the world to be worthless, like counting blades of grass? Let me say a bit about each of these cases in turn. First, consider the immoral end. Taking others' ends as reason giving involves recognizing them as having a kind of authority. This authority is bounded by morality. The choice of an immoral end gives one no reasons at all not to interfere with that end let alone reasons to aid the agent who chose the end. Notice that there may nonetheless be constraints on how we respond to immoral ends. People's rights can generate duties not to interfere in certain ways with even immoral actions. Suppose that Todd is hurling insults at Karl. He has the end of making him feel badly about himself. Todd's right to his body means that I am not permitted to forcibly restrain him by, for example, gagging him. This does not mean that I take Todd's end to be at all reason giving but rather that his rights place limits on how I may interact with him even though I do not take his end to give me reasons.

Next, consider the seemingly worthless end. Eli is busy counting blades of grass. You can see absolutely no value in this activity.<sup>10</sup> Let us set aside the possibility that Eli is here squandering his talents and so violating a duty to himself—that would push the case into the former category. Let us also suppose that Eli is not in the grip of some kind of compulsive behavior—that would raise questions about whether he is setting ends at all. On a certain kind of constructivism about reasons, it would seem that the mere fact that Eli has autonomously chosen to count blades of grass actually gives that activity value. So, on this kind of view, even if the activity looks worthless to you, Eli's choice makes that assessment a mistake. But I do not think we need to accept this kind of constructivism in order to arrive at the conclusion that Eli's end is reason giving. Consider a more realist picture on which there are reasons in favor of or against counting blades of grass that are independent of Eli's choices. Eli might then choose to engage in an activity that he really has no good reason to engage in. But even in such a case, it still seems that it would be disrespectful to walk by and make no effort at all to avoid disturbing Eli's count. It seems that respecting Eli requires respecting his choices, worthless though they may be. There are many ways a realist might motivate accepting this conclusion without claiming that Eli's choice genuinely imbues counting blades of grass with value. It might be, for example, that a kind of epistemological humility requires giving Eli the benefit of the doubt and so taking his end to be reason giving even though one sees no value in it.

Much more, of course, would need to be said to fully flesh out the grounds of the duty of beneficence and its contours. My aim in this section has just been to give a brief outline. In the remainder of Section 1, I turn to arguing that our duty against deception should be understood as an imperfect duty with a structure analogous to that of the duty of beneficence.

### *1.2. Agency and True Belief*

We are believers. That is, we all have countless beliefs and we are constantly forming new beliefs, both through deliberation and through other processes. Beliefs bear a significant relationship to truth, though there is considerable disagreement about how to characterize this relationship. Perhaps truth is the constitutive aim of belief. Or perhaps truth is

---

<sup>10</sup> John Rawls raises this case in *A Theory of Justice: Revised Edition* (Cambridge, MA: The Belknap Press of Harvard University Press, 1999) 379-380.

the standard of correctness for belief. Maybe yet another characterization best captures the relationship between truth and belief. For my purposes here, we need not settle this.<sup>11</sup> I will be arguing that, regardless of how belief itself is related to truth, true belief bears an important relationship to agency.

The relationship between having true beliefs and agency is obscured if we focus on only the relationship between our pursuit of particular ends and having true beliefs. There are costs to forming true beliefs. There are, of course, the costs in time, energy and other resources that are needed to form true beliefs. There is also the possibility that the truth itself may be disturbing and hence that we will suffer for believing it. Finally, although we might plausibly think that having true beliefs will generally enable us to pursue our particular ends better, there may be situations in which our ends would be more likely to be realized if we had false beliefs. Consider the aspiring pianist who will only have the confidence and motivation needed to be successful eventually if he believes falsely that he is doing well now. So, it is an open question whether having true beliefs will further our pursuit of any particular end.

But there are two noteworthy ways in which having true beliefs is more robustly connected with our agency. First, notice that our beliefs shape not only our pursuit of the particular ends we have chosen but also define the space of possible ends. Consider a mundane example. You believe that the new noodle shop on the corner has not yet opened. So, you do not consider going there for dinner. In this way, your beliefs shape the ends you take to be eligible for adoption. Our beliefs thus place a constraint on our agency. This gives us reason to be concerned that this constraint is tracking the truth.

The second way in which true beliefs are intimately connected with agency will require a bit more set up to appreciate. Let us begin by observing that successfully guiding oneself requires that the outcome at which one aims comes about in roughly the way one envisions. So, suppose you are playing golf. You hit the ball in a way that would quite likely have resulted in a hole-in-one. But a falcon swoops down and grabs the ball in flight. When the falcon realizes the ball is not prey, she drops the ball and, as it happens, it falls directly in the hole. Here the outcome you were aiming at is the outcome that obtains. Yet, there is a meaningful sense in which the outcome is not your achievement. Regardless of how we decide to score the 'hole-in-one', you cannot take credit for it in the way that you would have had the outcome been brought about in the way you envisioned when you hit the ball. And this is the sense in which it is a necessary (though not sufficient) condition on successful self-guidance that the outcome at which one aims come about in roughly the way one envisioned.

This means that successful self-guidance requires that our understanding of what we are doing and what is actually happening line up. Consider another example. You are trying to get to Pittsburgh but you mistakenly get on a plane headed to Philadelphia.<sup>12</sup> On

---

<sup>11</sup> Different versions of the relationship between belief and truth have been defended by many people. For some examples, see David Velleman, "On the aim of belief" in David Velleman (ed.), *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000); Nishi Shah, "How Truth Governs Belief," *The Philosophical Review*, 112:4, (2003) 447–82; Nishi Shah and J. David Velleman, "Doxastic Deliberation," *The Philosophical Review*, 114:4, (2005) 497–534; Pamela Hieronymi, "Controlling Attitudes," *Pacific Philosophical Quarterly* 87 (2006) 45–74; David J. Owens, "Does belief have an aim?" *Philosophical Studies* 115 (3): (2003) 283–305; Asbjørn Steglich-Petersen, "Weighing the aim of belief," *Philosophical Studies*, Volume 145, Issue 3, (2009) 395–405.

<sup>12</sup> I owe this example to Grace Helton and I am indebted to her for a helpful discussion of it.

the way to Philadelphia, the plane encounters bad weather and is rerouted to Pittsburgh. In this case, you aimed to go to Pittsburgh and you ended up in Pittsburgh. But you did not get there in the way envisioned when you embarked on your journey. Your understanding of what you were doing and what was actually happening did not line up. So, it is a happy coincidence that the outcome you were aiming at is the one that obtained. In this way, failures to understand how what you are doing is related to what is happening result in a defect in your attempt to guide yourself. Importantly, this defect in self-guidance may not reflect any defect in your reasoning about your situation. Your evidence, for example, might have decisively supported the false belief that you were boarding the Pittsburgh flight. So, in describing this situation as involving a defect in self-guidance, I do not mean to suggest that better choices were rationally open to you. Instead, I mean to indicate that there is a meaningful sense in which arriving in Pittsburgh is not your achievement.

I am not going to say much more about the nature of the defect in self-guidance I have just described. There are a number of qualifiers in light of which one might assess actions: autonomous, intentional, full-blooded, and so on. I am making no claim about which, if any, of these qualifiers best captures why failing to understand what one is doing and how that is related to what is happening poses a problem for agency. These qualifiers mean different things in the hands of different theorists and wading through these distinctions would take us too far afield. For my purposes, it is sufficient to observe that something has gone wrong in the kinds of cases I have described.<sup>13</sup> To remain neutral between these various qualifiers, I will stick with the description I gave above: not understanding what one is doing and how that is related to what is happening results in a defect in one's attempt to guide oneself.

What is involved in having this kind of understanding? For the most part, we can characterize understanding one's actions and their context in terms of having true beliefs about them. But there are two important qualifications here. First, one might think that in certain domains truth is not the right standard of success. For example, a fictionalist about numbers might hold that claims about numbers are not literally true. Yet accepting certain claims about numbers may plausibly be part of understanding one's actions and their context. If one adopts these kinds of views, one might replace my talk of truth with a more capacious standard of success.

Second, one might wonder whether having true beliefs is sufficient for understanding one's actions and their context. If one has true but unjustified beliefs there is a sense in which one's success is again a happy accident. Indeed, Gettier cases suggest that even relying on justified true beliefs can involve a kind of luckiness. So, perhaps instead of thinking that true beliefs suffice for understanding one's actions and their context, we should instead think that something more is required.

---

<sup>13</sup> This kind of defect might seem foreign to the Kantian framework in which I am operating. Kant is most famous for identifying autonomy with following the Categorical Imperative. And there is no reason to think that mistakenly getting on the plane to Philadelphia involves willing in a way that is contrary to the Categorical Imperative. But the Kantian view has room for evaluating actions along more dimensions. Kant himself distinguishes autonomy, i.e. internal freedom, from independence, i.e. external freedom. While autonomy is characterized in terms of bearing the right relationship to one's own inclinations, independence is characterized in terms of bearing the right relationship to other agents. One might think that the defect I have identified in the text suggests that actions may also be evaluated in terms of their relationship to the world.

Two observations are important in response to this worry. First, it is in principle possible to distinguish between epistemic and agential luckiness. If you form a true but unjustified belief that the plane you are boarding is headed to Pittsburgh, there is no mismatch between what you take yourself to be doing and what is actually happening. Here you are lucky that what you believe happens to be true, but the choices you are making are engaging with the way the world actually is and hence arriving in Pittsburgh seems to reflect your self-guidance.

Secondly, though, this does not mean having justification or knowledge is irrelevant to one's agency. These states enable one to guide oneself successfully in a wider range of circumstances than the present one. Mere true belief, as Socrates informs Meno, is apt to 'run away'. So, we have an agential interest in something more than true belief, but the nature of this agential interest differs. For the most part, I will be concerned with cases in which only our agential interests in true beliefs are in play because they are cases in which one agent tries to give another a false but justified belief. But the account will have room for considering the practical relevance of justification in more complex cases.

So, setting aside the complications just canvassed, I maintain that we have an agential interest in the truth. False beliefs constrain our agency by limiting the space of possible ends and undermine successful self-guidance by preventing us from understanding our actions and their context. Three important clarifications are in order. First, the focus on agency in the account I am giving reflects the broadly Kantian outlook that I am working within. This marks an important contrast with other recent attempts to explain the significance of true belief by connecting true belief with happiness.<sup>14</sup> But I have not attempted to argue for the fundamental importance of considerations about agency in our ethical outlook. That is something I am taking as a starting point. Secondly, though, I have not claimed that well functioning agency ought to matter more to us than anything else. In Section 1.5, I will consider cases in which agency may matter less than something else and hence cases in which the truth of our beliefs may reasonably matter less to us than something else. So, Section 1.5 will give some more content to the kind of significance my view attributes to agency. Third, I am focusing on a normative rather than a descriptive claim. I claim that the relationship between true belief and our practical agency makes the truth of our beliefs an object of significant practical concern. This is not a claim about how people in fact regard the truth. In Section 1.5, I will also consider cases in which people do not treat the truth as having the practical significance it in fact has. For now, however, I want to simply focus on the normative claim.

---

<sup>14</sup> For accounts that connect the significance of true belief with happiness, see, for example, Michael P. Lynch, *True to Life: Why Truth Matters* (Cambridge, Mass: MIT Press, 2004) and Linda Zagzebski, "The Search for the Source of Epistemic Good," *Metaphilosophy* 34:1/2, pp. 12-18. For objections to such views, see Allan Hazlett, *A Luxury of Understanding: On the Value of True Belief* (Oxford: Oxford University Press, 2013). Hazlett's objections do not apply straightforwardly to linking the significance of true belief with agency in the way I have suggested. Hazlett briefly considers a view that makes such a connection and writes: "Not knowing the right word for 'sixteen' in French is bad because someone ignorant of this runs the risk of bad consequences, but there is nothing bad, per se, about actions (or 'actions') based on false beliefs. There is nothing eudaimonically valuable about performing actions, per se, or about performing fully autonomous actions, fully intentional actions, and so on." (81). I suspect that there are plausible conceptions of eudaimonia that would not affirm Hazlett's claim here. But in any case, my account does not attempt to draw any necessary connection between eudaimonia and agency. Hazlett does not consider the possibility that agency might be significant in its own right, and that is the claim on which the Kantian outlook relies.

### 1.3. *The Duty of Doxastic Concern*

The agential interest we have in the truth grounds a duty with respect to others' beliefs that is analogous to the duty of beneficence. Recall that the duty of beneficence requires us to treat the ends of others as reason giving. This is a way of showing respect for others as agents by respecting exercises of their agency. The agential interest we have in the truth is not something we choose in the way we choose ends. Rather, it is an implication of the structure of our agency. I suggest that respecting others as agents requires taking this agential interest in the truth as reason giving in much the same way that we take others' ends as reason giving. That is, we have a duty to have others' agential interest in the truth as an end. Let us call the duty that reflects this requirement the duty of doxastic concern.<sup>15</sup>

Using the duty of beneficence as a model, then, it seems that we have reasons to help others form true beliefs and reasons not to interfere with their formation of true beliefs. These reasons are distinct from the reasons that the duty of beneficence itself gives us with respect to others' beliefs. For example, suppose someone asks you for directions. Taking her ends as reason giving means that you have reason to give her this information as a means to her end of arriving at her destination. This reason may be stronger or weaker depending on what her destination is, e.g. the hospital or the theater. Whether the duty of beneficence requires stopping to give directions depends, as I have indicated, on the other ends in play.

The reasons we have stemming from the agential interest others have in the truth will often work to augment the reasons we have stemming from their particular ends. So, even if I do not have a duty to give you bus fare, I may have a duty to give you directions to the bus stop—even supposing that the change and the time are worth the same to me. Although your end may not be significant enough to give me a duty to help you, your agential interest in the truth might be.

Additionally, the agential interest others have in the truth gives us reason to give them information that may not be a means to any particular end that they have chosen but instead gives them reason to adopt new ends or reconsider the ones they have chosen. If, for example, you are planning a trip to Italy, I have reason to tell you about the excellent gelato. This information may not contribute to any particular end you had already adopted but contributes to your understanding of the space of possible ends.

Finally, notice that the reasons we have against interfering with others' formation of true beliefs are generally stronger than the reasons we have to help them do so.<sup>16</sup> To see why, consider once again the model of the duty of beneficence. If I take your end as reason giving, I must have some sufficient competing reason not to help you realize it. But to work against the realization of your end requires even stronger competing reasons. This is why it is easier to justify declining to help the tourists take their photo than it is

---

<sup>15</sup> The reasons we get from others' agential interest in the truth may sometimes conflict with one another. I may, for example, need to tell you a falsehood to get you to believe a truth. Or it may be that telling you a truth will lead you to believe a falsehood (as, for example, in a 'double bluff'). In these cases, I cannot lead you to an entirely veridical set of beliefs. What the duty of doxastic concern directs me to do in this case will then depend on other elements of the case, like the relative importance of the beliefs. Of course, if there is a distinct duty not to lie, that will complicate what I ought to do.

<sup>16</sup> As I go on to explain in Section 1.4, I do not mean to suggest here that there is any deep, metaphysical distinction between the actions that count as helping and those that count as interfering. The distinction will often rest on people's conventional expectations.

to justify running through their shot. This same line of reasoning applies to the reasons we have stemming from the agential interest in the truth. And notice that in this case interference will typically have an even deeper impact on the other's agency. The tourists whose photo was disrupted can assess how to move forward. In contrast, interfering with a person's understanding of the world leaves her poorly situated to understand the space of possible ends and so to direct herself. This suggests that, other things being equal, the reasons we have not to interfere with others' formation of true beliefs are even stronger than the reasons we have not to interfere with their particular ends.

Let us consider how the duty of doxastic concern plays out in concrete cases. First, consider the case of Jill and George, who are neighbors and acquaintances. George is a serious basketball fan and often starts long conversations about basketball when they both happen to be doing their laundry, taking out the trash, etc. Tonight Jill is coming home late from a basketball game. Jill may be able to help George form a wide variety of true beliefs, say, beliefs about the orbit of the planets, the mating habits of penguins and the history of the paperclip. George's agential interest in the truth gives her some reason to discuss any of these with him. There are two considerations, however, that suggest that Jill needs very little reason to decline to help on these matters. First, Jill has no particular reason to think that George's choices are being influenced by his beliefs on these matters. So, although George has an agential interest in the truth regarding these matters, that agential interest is quite minimal. Second, Jill has no particular reason to think that George has false beliefs regarding these matters. George may be quite aware that he knows nothing about the orbit of the planets. In contrast, if George believed, say, that the planets orbit around Earth, Jill's reason to help George would be somewhat stronger. Although the orbit of the planets is of minimal practical concern for George at present, in this case George would be poorly situated to respond if this became relevant to his choices in the future. Still, Jill's reason to help George on this topic remains easily outweighed.

But now consider the reason Jill has to help George form true beliefs about her attendance at the basketball game. She has reason to think that George's agential interest in the truth about this is more substantial—believing the truth would affect his choices. This gives her somewhat stronger reason to help George. This reason might still be outweighed. If she needs to get up early in the morning for a meeting, she might decline to help George, not mentioning the details of her evening when she passes him in the hall. And this seems like the right result—omitting this information in this case is not morally problematic.

Now suppose that Jill has reason to think that her silence will lead George to form a false belief. George will assume that she is coming home from late from work, something she often does. Here Jill has somewhat stronger reason to help George. George not only lacks the truth about something that would affect his choices, he is poorly situated to respond to this lack. Nonetheless, it still seems that Jill's reasons to help may be outweighed by her early morning meeting. And again, this seems like the right diagnosis. Although Jill has stronger reason to help, she may still decline to do so in light of her competing reasons.

Now suppose that Jill does not merely omit information but in fact behaves in a way designed to lead George to form false beliefs about her evening activities. Before George spots her, she pulls a stack of files out of her satchel to make it appear that she is coming home from working late. My view suggests that this action stands in need of more

justification than the omission. Jill is working against George's agential interest in the truth. She needs to have competing reasons that are sufficient to justify this. But my view suggests that there are circumstances in which working against George's agential interest in the truth may be permissible. Merely needing to get up early in the morning does not seem to be sufficient to justify this kind of interference. So we will need to imagine weightier competing reasons in place of this. In Section 1.5, I will describe ends that might more plausibly be taken to outweigh another's agential interest in the truth. But for now what is important is simply to notice the plausibility of the diagnosis being offered of this case: this kind of interference may be impermissible, but need not always be so. With this gloss on the practical implications of the duty of doxastic concern in mind, we can turn to the relationship between this duty and the duty not to deceive.

#### 1.4. *The Duty Not to Deceive*

In the foregoing discussion of Jill and George, I used the duty of doxastic concern to evaluate several actions, some of which might be considered deceptive. The final case in which Jill pulls out her satchel seems most clearly to fall in this category. Here Jill aims to lead George to form a false belief by exposing him to evidence for that belief. I have suggested that there are strong moral reasons against doing this but that these kinds of actions may be permissible depending on what else is at stake. The reasoning that led to this conclusion captures what seems intuitively wrong about deception. Deception interferes with a particularly deep aspect of the victim's agency. So, it should not be undertaken lightly.

We also considered two cases of omission. In the first, Jill could expand the range of practically relevant possible ends open to George by telling him about her evening. But here the omission does not lead George to any false beliefs—he will simply not form beliefs about how she spent her evening. In the second case, George will actually form a false belief about how Jill spent her evening. There has been disagreement about whether Jill counts as deceiving George in either of these cases.<sup>17</sup> But notice that we do not need to settle that disagreement in order to evaluate the moral standing of Jill's actions. Reflecting on the duty of doxastic concern suggests that Jill has an easily outweighed reason to help George in the first case and a somewhat stronger but still easily outweighed reason to help George in the second case. This is a very intuitive analysis of these cases.

If we have a very capacious understanding of deception and include even actions like those in the first case of omission, the duty against deception will just be the duty of doxastic concern as I have described it. I suspect, though, that most people would resist such a capacious conception of deception. Given a more limited conception of deception, the duty against deception will be one component of the broader duty of doxastic concern. For example, perhaps the duty against deception should be understood as reflecting just the reasons we have not to interfere with others' formation of true beliefs and leaving aside the reasons we have to help them form true beliefs. Note that distinguishing between these two sets of reasons does not require that there be any deep, metaphysical distinction between the actions that count as helping and those that count as interfering.

---

<sup>17</sup> For a view that describes both cases of omissions as cases of deception, see Roderick M. Chisholm and Thomas D. Feehan, "The intent to deceive," *Journal of Philosophy*, 74 (1977): 143–159.

The distinction will often rest on people's conventional expectations. But the distinction might still be useful in our moral reasoning as a way of highlighting how particular ways of acting must meet a very high justificatory burden in order to be compatible with the concern we ought to have for others' doxastic situations. In principle, though, we do not need the category of deception to appreciate the general structure of our duty not to engage in such actions without meeting this justificatory burden. It turns out, then, that we do not need to settle on a particular conception of deception in order to appreciate the moral status of the actions that are typically described as deceptive or even to settle the moral status of borderline cases.

### *1.5. The Agential Interest in the Truth Reconsidered*

With this account of the duty of doxastic concern in mind, we can turn to considering when it is permissible to act contrary to the reasons we have stemming from others' agential interest in the truth. As I noted at the outset, I have left open the possibility of cases in which the truth ought to matter less to us than something else. There are two categories of cases I will explore in connection with this possibility. First, I will consider the possibility of needing to choose between addressing one person's agential interest in the truth and preserving another's agency, e.g. cases like the murderer at the door. Second, I will consider the possibility of needing to choose between addressing a person's agential interest in the truth and preserving that very person's agency, e.g. cases in which paternalistic interference might seem justifiable. I have also acknowledged that people may not always treat the truth as having the practical significance it in fact has. I will end by considering the appropriate response to such cases.

To begin, notice that addressing a person's agential interest in the truth may sometimes be at odds with responding to the serious agential interests of others. The case of the murderer at the door is a classic example. The murderer at the door wants to know the whereabouts of his intended victim. But of course giving the murderer that information will put his victim's life at risk. The framework I have suggested can give an intuitive diagnosis of why the duty of doxastic concern permits deceiving the murderer. Of course, if there is a distinct duty not to lie, we will need a separate explanation of why lying to the murderer at the door is permissible. I will return to this issue Section 2.2. For now, let us set aside any distinct objection to lying and consider how all the other reasons in this case stack up.

First, notice that the murderer's immoral end gives rise to no reasons of beneficence at all to help him. So, the reasons to help the murderer are limited to his agential interest in the truth. Although this is a very significant source of reasons, we have already seen that these reasons are to be compared with the other reasons in play. And here the victim's life gives rise to reasons that are much more significant. So, of course, we have no duty to be forthcoming with the murderer. And these very same reasons suggest that we are permitted to interfere with the murderer's agential interest in the truth for the sake of his victim's life.

Next, notice that addressing a person's agential interest in the truth can sometimes also be at odds with preserving that very person's agency. Consider Sarah, a recovering drug addict. She was devastated by her recent break up. You have good reason to believe that if she learns that her ex is already with someone new, she will return to drug abuse and this will leave her ability to direct herself severely impaired.

I think this is a genuinely difficult case and that the framework I have outlined reveals why. Our reasons for not withholding the information are tied to Sarah's agential interest in the truth. But giving her that information threatens another deeply important aspect of her agency.

Nonetheless, the framework suggests a natural direction for addressing this situation. Sarah's agency is at risk and so, to the extent possible, Sarah should be the one to decide which risks to take. So, suppose she asks you if her ex is with someone new. The thing to do if possible is to let her know you think that it is dangerous for her to investigate that matter. If she persists, you may decline to help her. She is entitled to take what risks she thinks best, but her agential interest in the truth does not compel you to aid her in this case. But this does not ground a permission to deceive her. Deceiving her would indicate that you think that you rather than she should decide what risks she should take. And that would fail to respect her agency.

Finally, let us consider cases in which people prefer to remain ignorant on certain matters or even to be deceived, not just for the sake of preserving their agency, but for the sake of particular ends. This need not be a mistake, but it can be. Let me begin with the unproblematic cases. As cognitively limited agents, we are not able to know everything. So, there are many things one might prefer to remain ignorant about so as to use one's limited cognitive resources well. This consideration alone, however, does not generally support deception.

A preference for being deceived often stems from concern with how the truth will affect one's ability to pursue some specific end. Suppose that Eric tells his fiancée Ann that if she ever cheats on him she should deceive him about that.<sup>18</sup> This is because Eric is committed to staying married to Ann come what may, but worries that he will have a hard time living with infidelity. There are, I think, two attitudes that could motivate Eric's preference. On the one hand, Eric might think that his marriage is so important that it is worth sacrificing his agency. Since there are ends that one could reasonably give up one's life for, it seems clear that there are ends one could reasonably sacrifice some agency for.<sup>19</sup> On the other hand, Eric may simply fail to take seriously the loss of agency involved in believing falsehoods. In this case, Eric fails to recognize the agential interest in the truth he has in virtue of the structure of his agency. And in this way he fails to properly respect himself.

Given that the best case for deceiving Eric is one that involves the sacrifice of his agency, it seems clear that Eric is entitled to decide for himself whether this is a sacrifice he wants to make. Ann needs his consent. So, it matters greatly that Eric is asking Ann to deceive him. Even if Ann thought he ought to be willing to make this sacrifice, she is not entitled to impose it on him. So, in these kinds of cases, the permission to deceive importantly depends on whether the person has indicated a willingness to be deceived. And in these cases the deception is perhaps not best described as paternalistic – the deceiver simply abides by her interlocutor's views about the relative importance of his agency.

---

<sup>18</sup> Christine M. Korsgaard calls into question whether one can really agree to be deceived. See Korsgaard (1996) 155-156. But Korsgaard's incredulity does not seem well founded in light of such commonplace phenomena as setting the clock forward a few minutes in order to trick oneself into being on time. For further discussion of this issue, see Jonathan E. Adler, "Lying, Deceiving, or Falsely Implicating," *The Journal of Philosophy*, 94:9 (1997) 440-441.

<sup>19</sup> This claim marks an important difference between the view I am describing and some other Kantian views. See, for example, J. David Velleman, "A Right of Self-Termination?" *Ethics*, Vol. 109, No. 3 (Apr., 1999) 606-628.

Even if Ann is permitted to deceive Eric, it still seems that she can decline to do so. In choosing to be deceived, Eric is choosing not to guide himself but instead to be guided by Ann. This is a responsibility Ann might reasonably choose to decline. And if the case involves a lack of self-respect instead of self-sacrifice, Ann should decline. She should not facilitate Eric's failure to respect himself. In this way, Eric's agential interest in the truth continues to shape the reasons Ann has regardless of Eric's own attitudes towards the truth. So, I think this account can give a plausible diagnosis of our duties in both cases in which the truth matters less than something else and in cases in which an agent mistakenly believes this to be so.

### *1.6. Complementing the Duty of Doxastic Concern*

Before turning to whether there might be a distinct duty not to lie, I want to highlight two other kinds of duties that may also govern how our actions convey information to others. These duties are an important complement to the duty of doxastic concern. To begin, let us contrast the case of Jill and George with the case of John and his wife Carol. John has come home late after having drinks with his friends. He believes that knowing this would affect Carol's choices—she would want to discuss it. But John has an early meeting in the morning. Suppose he simply fails to mention his evening activities, thinking that Carol will assume that he is coming home late from work. Here even the omission seems morally problematic, suggesting that John has stronger reason to share the information with Carol than Jill does with George. In this way, certain kinds of relationships, like the relationship between spouses, seem to affect the reasons we have to help. Examining exactly how and why this is would take us too far afield. Here I just want to draw attention to this familiar component of our everyday experience. If John fails to tell Carol about his evening, he may not be failing to abide by the duty of doxastic concern that he owes all people but instead failing to abide by the duty of disclosure he has to his spouse. In this way, we can choose to enter relationships that give us distinct reasons to attend to the doxastic situation of those with whom we have those relationships.

Next, consider a second version of the case of John and Carol. Suppose that while Carol is out running errands, she actually spots John at the pub with his friends. Their eyes meet and so his evening activities are common knowledge between them. When he gets home, he acts as if nothing has happened and goes through the motions of coming home late from work. Here John has no reason as a spouse or as a fellow person to attend to Carol's doxastic situation. He knows that she knows the truth. But notice that giving Carol new information is not the only reason he may have to tell her about his evening. Whether he acknowledges what happened or not may affect the character of their interactions going forward. For example, leaving it to Carol to bring up what she saw may either be a way of empowering her or silencing her, depending on the larger dynamics of their relationship. Thus, we may have reasons to be forthcoming (or not) that are orthogonal to those that the duty of doxastic concern is tracking.

## **2. The Duty not to Lie**

I have argued that we have a duty of doxastic concern that is analogous to the duty of beneficence. Depending on how we define deception, the duty of doxastic concern either

just is the duty not to deceive or is a broader duty that encompasses the duty not to deceive. Either way, reflecting on the duty of doxastic concern leads us to plausible moral evaluations of the kinds of actions we typically describe as deceptive, at least insofar as they do not involve communication.

Suppose that, rather than pulling files out of her satchel, Jill simply tells George that she was working late when in fact she was at the basketball game. This is a paradigmatic case of a lie. Lies will typically be interferences with others' possession of true beliefs.<sup>20</sup> Insofar as this is the case, the duty of doxastic concern requires that they meet a high justificatory burden. And this will also be true for other types of speech that interfere with others' possession of true beliefs.

Is there a duty not to lie that is distinct from the duty of doxastic concern? I will not be able to fully settle this question here. Instead, here I aim to consider how the duty not to lie, if there is one, would function in our moral reasoning. I will use an analogy with the perfect duty to keep one's promises to show how such a duty would operate. This analogy will demonstrate how a duty to perform or refrain from particular actions rather than to have a certain end can be a plausible part of our moral reasoning.

Two points of clarification are in order. First, the proposal does not claim that lying literally involves breaking a promise but only that the duty not to lie has the same structure as the duty to keep one's promises. While others have drawn a tighter connection between lying and promise breaking, my proposal is limited to identifying illuminating similarities.<sup>21</sup> Indeed, we will see that while these duties have important similarities, they also have important differences. Second, I will begin by focusing on paradigmatic cases of lying without relying on any particular definition of lying. But in the course of the discussion we will have occasion to consider how questions about the definition of lying bear on our moral reasoning.

The discussion will proceed as follows. First, I will examine the idea that lying is in some sense worse than other forms of deception and argue that my proposal suggests a different and more plausible way of construing the relationship between lying and other forms of deception. Next, I will consider a potentially powerful objection, namely that a perfect duty not to lie would be implausibly strong. I will argue that thinking through the example of promises suggests a way of addressing this objection. Finally, I will examine how the debate about the definition of lying bears on our moral reasoning.

### 2.1. *The Comparative Claim*

At the outset of this paper, I suggested that people often take the duty not to lie to be more stringent than the duty not to deceive. This is not the same as the claim that lying

---

<sup>20</sup> But lies need not interfere with others' possession of true beliefs. Consider a helpful example of a 'bald-faced' lie from Thomas L. Carson. A student might deny cheating on a test for procedural reasons even though he knows that the dean to whom he lies will not be deceived. See Carson (2006) 290. See also, Roy Sorensen, "Bald-faced lies! Lying without the intent to deceive," *Pacific Philosophical Quarterly* 88:2 (2007) 251-264. The case of a lie that is not intended to deceive is much like a promise that it does not benefit the promisee to keep, like a promise to someone who hopes you will break the promise so that others will come to distrust you. Although these are not the standard cases of lying or promising, these cases help draw out how there might be distinctive duties not to lie and to keep one's promises that do not just reduce to the duty of doxastic concern and the duty of beneficence respectively.

<sup>21</sup> See, for example, W. D. Ross, *The Right and the Good*, (Oxford: Oxford University Press, 1930) 21. For a related view and a helpful discussion, see also Carson (2006).

is some sense morally worse than other forms of deception. I am going to argue that understanding the duty not to lie on the model of the perfect duty to keep one's promises supports the former but not necessarily the latter claim. Moreover, I will argue that the former claim actually better captures the intuitive difference between the moral status of lying and other forms of deception.

Jennifer Saul proposes the following counterexample to the claim that lying is worse than other forms of deception, like deliberately misleading:

Charla is HIV positive, but she does not yet have AIDS, and she knows both of these facts. Dave is about to have sex with Charla for the first time, and cautiously but imprecisely, he asks (3).

(3) Do you have AIDS?

Charla responds with (4).

(4) No, I don't have AIDS.

Charla and Dave have unprotected sex. Dave becomes infected with HIV. It is unquestionably true that Charla deceived Dave about her HIV status, and also unquestionably true that she did not lie – she merely mislead him. Yet it seems completely absurd to suppose that Charla's deception was even a tiny bit better due to her avoidance of lying.<sup>22</sup>

Consider the following potential analysis of this case using elements of the account I have been developing. The duty of doxastic concern indicates that Charla has a reason to be both honest and forthcoming with Dave. And given the stakes in this case, it is quite plausible that these reasons are decisive. That is, although Charla generally has discretion about whether and how to help Dave understand the context in which he is acting, in this case the only way to respect him is to tell him that she is HIV positive.

But now suppose that, in addition to the imperfect duty of doxastic concern, we also have a perfect duty not to lie. My suggestion is that distinguishing perfect and imperfect duties opens up the possibility of identifying an important difference between lying and other methods of deception without being committed to the claim that there is some sense in which using one method is always worse than using another.

To see what this might involve, I will begin with an analogy. As I indicated above, perfect duties are duties to perform or refrain from particular actions while imperfect duties are, in the first instance, duties to adopt a particular end. I have already described the imperfect duty of beneficence. Let us now contrast this with the duty to keep one's promises. This is a duty to perform certain actions rather than to have a certain end and thus a perfect duty. With this in mind, consider the case of Amelia and Patrick. Patrick is out sailing when he slips and falls out of his boat. He panics and so has difficulty swimming. Amelia is sailing nearby and sees him in distress. I have suggested that although the duty of beneficence is an imperfect duty, in this case it requires her to throw him a life preserver. But now suppose that earlier Amelia had promised Patrick that she would help if he encountered any problems while they were sailing. Now Amelia has a distinct reason to help Patrick, a reason that stems from her perfect duty to keep her

---

<sup>22</sup> Saul (2012), 72.

promises. Is it worse for Amelia to fail to throw the life preserver when she promised she would help? I think that there are plausible pulls in both directions on this question. But for my purposes here, we need not answer it. What is important is that we have identified a difference between Amelia's duties in the two cases that is in some ways orthogonal to the question of which failure is worse.

With this in mind, let us return to Charla's claim to Dave that she does not have AIDS. The imperfect duty of doxastic concern condemns this action. If we suppose she also has a perfect duty not to lie, her lie would be subject to a distinct criticism. The presence of a distinct criticism may or may not make her action morally worse. But it can potentially explain why there can sometimes be cases in which one is permitted to deceive using one method but not another. Just as there are cases in which one is permitted not to help another unless one has promised, there may be cases in which one is permitted to deceive another unless the deception takes the form of a lie. The analogy with promising thus opens up a way of making sense of how there could be a distinct duty not to lie. Such a duty would be directed toward a certain form of action rather than the effect of that action on the other party's doxastic situation.

## 2.2. *Are Perfect Duties Implausibly Strong?*

Since a perfect duty is a duty to perform or refrain from a particular action, it may seem as though describing the duty not to lie as a perfect duty saddles us with the conclusion that it is wrong to lie to the murderer at the door. In what follows, I consider how a perfect duty to keep one's promises can avoid a similar objection and then extend this reasoning to the duty not to lie.

I am going to examine two features of promising. First, valid or binding promises can be made only under certain circumstances. Second, there are limits on the extent of promissory obligations. Let us consider each of these features of promising in turn.

Promising is an exercise of what we may call a normative discretionary power. Normative discretionary powers enable people to change the permissions and obligations that were in place prior to their exercise. Consenting is another familiar exercise of a normative discretionary power. Whereas consenting overrides a barrier to the permissibility of an action, promising generates an obligation to perform the promised action.

Wrongful coercion prevents the successful exercise of normative discretionary powers.<sup>23</sup> To see why, let us begin with a case in which wrongful coercion limits one's ability to consent. When a mugger says, "You're money or your life," he is thereby impermissibly constraining your options. This makes it the case that your action cannot count as a reflection of your discretionary authority with respect to your money – you lack the options you are entitled to have discretion with respect to. So, the act of handing the money over to the mugger does not have its normal normative significance – it does not amount to consenting to let the mugger keep the money. Likewise, a wrongfully coerced promise does not have its normal normative significance. A promise extracted at gunpoint to bring the mugger money tomorrow is not binding. So, even if one has a perfect duty to keep one's promises, one need not bring the mugger the money.

---

<sup>23</sup> See Japa Pallikkathayil, "The Possibility of Choice: Three Accounts of the Problem with Coercion," *Philosophers' Imprint*, Vol. 11, No. 16 (November 2011): 1-20.

Next, let us consider the limits on the extent of promissory obligations. If I promise to meet you for lunch at noon but find that I must stop to help a drowning man, I ought to do so even if that makes me late. To understand why, first notice that one may add conditions to what one promises. I may, for example, promise to meet you for lunch at noon unless my car breaks down. But it would odd and seemingly unnecessary to say that I promise to meet you at noon unless I have to stop to help a drowning man. Morality places conditions like this one within all promises. A full exploration of these conditions would take us too far afield. Here it suffices to observe that taking the extent of promissory obligations to be limited by people's serious agential interests is very intuitive.<sup>24</sup>

Moreover, this proposal is consistent with continuing to regard the duty to keep one's promises as a perfect duty. Perfect and imperfect duties play different roles in our moral reasoning. Imperfect duties are duties to have certain ends. Having these ends gives one reasons to act in certain ways, but only gives one decisive reasons in certain circumstances. The reasoning associated with an imperfect duty takes the form of weighing the reasons associated with the required end against other reasons stemming from other ends.

In contrast, perfect duties are duties to perform or refrain from particular actions. As a result, perfect duties do not have the weighing structure of imperfect duties. Instead perfect duties preempt weighing countervailing reasons.<sup>25</sup> So, if I promise to help you wash your car, my promise is not just one reason among others that I compare with all the other reasons I have, like the reason I have to go to a concert. But this is consistent with taking morality to place limits on the extent of promissory obligations. This is just to say that promises preempt some but not all countervailing reasons.

So, we have before us two reasons why we may be permitted to act in ways that appear to be contrary to a promise. Sometimes promises are not binding because they have been given in problematic circumstances. And sometimes one may be permitted not to do what one has promised because of conditions limiting the extent of one's promissory obligation. With this in mind, let us turn to the duty not to lie.

I am going to argue that the duty not to lie is structurally similar to the duty to keep one's promises in significant respects. But it is important to note at the outset that there is one respect in which the duty not to lie is substantially more complicated. This is a duty not to engage in certain kinds of speech. But speech is governed by both linguistic and moral norms. In what follows, we will see that these two sets of norms are intimately intertwined in this duty, but still importantly distinct.

So, consider the following proposal. We have a perfect duty not to say things to people that we believe to be false.<sup>26</sup> Note two qualifications. First, I am using 'say' here in an informal sense as a placeholder. In what follows, I fill in that placeholder by focusing

---

<sup>24</sup> But if it turns out that I was not obligated to meet you at noon given the circumstances, why is it appropriate for me to apologize or try to make it up to you? Notice that even if I lack a promissory obligation to meet you under the circumstances, I am still disappointing your expectations by failing to be on time. And this may give me reasons to approach the situation in the aforementioned ways.

<sup>25</sup> For a helpful discussion of this form of reasoning, see Joseph Raz, "Promises and Obligations," in *Law, Morality, and Society: Essays in Honour of H.L.A. Hart* (Oxford: Oxford University Press, 1977) 210-228.

<sup>26</sup> It is, of course, often permissible to say something one believes to be false when engaged in ironic or metaphorical speech. In what follows, I leave open how best to block the implication that these forms of speech are impermissible. Maybe simply tacking on the stipulation that one is not engaged in these forms of speech to the description of the duty is the best way of handling this. Alternatively, filling out the duty in terms of assertion, as I go on to do, may place these forms of speech outside the scope of the duty.

on assertion, but the placeholder could in principle be filled in by another notion if that proved more keeping with what we are doing in our everyday lives when we talk to people.<sup>27</sup> Second, in this section I will provisionally treat the duty just described as equivalent to the perfect duty not to lie. I will consider debates about the definition of lying in Sec. 2.3 and examine this provisional definition of lying in that context.

So, suppose we have a perfect duty not to assert what we believe to be false. Different views of assertion propose different connections between asserting and changing a speaker's responsibilities. For example, perhaps asserting involves committing oneself to either defend or retract what is said in light of criticism.<sup>28</sup> Or perhaps asserting involves being in some sense responsible for the truth of what one asserts.<sup>29</sup> For my purposes, we need not sort through the different possibilities. It is enough that assertion has some responsibility shifting element or another.

In light of this element of assertion, it is plausible to think that wrongful coercion affects the possibility of successfully asserting in much the same way that it affects the possibility of successfully promising. Wrongfully coerced speech does not have the normal normative significance, whatever that might be.<sup>30</sup> And so a wrongfully coerced speech falls outside of the scope of the perfect duty described above because such speech does not constitute successful assertion.<sup>31</sup>

This is a helpful first step in explaining the permission to lie to the murderer at the door. A plausible way of filling out that case involves thinking of the murderer as implicitly threatening to harm you if you do not tell him what he wants to know. But there are other ways of telling the story where the murderer intends you no harm. He may, for example, mistakenly think that you will be glad to see his would-be victim dead. Moreover, there are cases of permissible lies in which there is no wrongdoing anywhere in the picture. You may, for example, need to lie to me in order to prevent a disaster because explaining the truth to me would take too long. In these cases, no wrongful coercion deprives your speech of its normal normative significance. So, what explains the permission to lie in such cases?

Here we can once again use the duty to keep one's promises as a model. I suggested above that a perfect duty can still be limited in application. It is plausible to think that the duty not to assert what one believes to be false has certain built-in limits much like the duty to keep one's promises. The key difference here is that there is another set of norms in the picture that may not have any such limitations, namely the linguistic norms. It may well be that in asserting what one believes to be false one is flouting linguistic norms. So, here we get a sense for the complex relationship between the linguistic and moral norms. The moral norms are in a way dependent on the linguistic norms because

---

<sup>27</sup> For example, perhaps we should be speaking instead of warranting. See Carson (2006) and Saul (2012).

<sup>28</sup> Robert Brandom "Asserting," *Noûs* 7 (1983) 637–650.

<sup>29</sup> Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000).

<sup>30</sup> For a helpful discussion of this point, see Katherine Hawley, "Coercion and Lies," forthcoming in *Lying: Language, Knowledge, Ethics, Politics*, edited by Eliot Michaelson and Andreas Stokke (Oxford University Press).

<sup>31</sup> Sorensen (2007) worries that this view cannot do justice to the extent to which people attempt to avoid saying certain things even when coerced, for example the coerced denial of one's religious beliefs. The discussion in Section 1.6 of how we can have reasons for and against speaking that are orthogonal to both the duty not to lie and the duty of doxastic concern can go some way towards addressing this worry. For further discussion of this issue, see Tim Kenyon, "Cynical Assertion: Convention, Pragmatics, and Saying 'Uncle,'" *American Philosophical Quarterly* 40:3 (2003) 241–248.

they reference phenomena that are defined in terms of the linguistic norms, e.g. assertion. But the moral norms are not beholden to mirroring the linguistic norms.<sup>32</sup>

I take the above discussion to show that a perfect duty not to lie need not be implausibly strong. Of course, this falls far short of showing that there is any such duty. I have made no attempt to show why we should think there is a perfect duty to not to assert what one believes to be false. Among other things, working out such an account would require settling on a view of assertion. And that is beyond what I can do here. Moreover, looking at the duty to keep one's promises is instructive regarding the prospects for defending the existence of a duty not to lie. While just about everyone agrees that there is a duty to keep one's promises, the grounds of that duty are a matter of significant debate. Determining the grounds of the duty not to lie is likely to raise many similarly difficult and contentious issues. So, an adequate exploration of these grounds goes beyond what I can do here.

### 2.3. *The Definition of Lying*

In the previous section, I provisionally defined lying as asserting what one believes to be false. There are many respects in which this definition might be contentious. For example, perhaps assertion is not the right category of speech to focus on.<sup>33</sup> I acknowledged this possibility in the more general description of the duty under consideration. Or perhaps instead of focusing on what one believes to be false, we should instead focus on what one does not believe to be true.<sup>34</sup> This would leave open the possibility of lying by saying something that one is agnostic about. This by no means exhausts the potential dimensions of disagreement. But just looking at these will be instructive.

The definitional debate is useful insofar as it brings to our attention different kinds of actions and allows us to inquire about the duties we have with respect to them. But we do not need to settle the definitional debate in order to proceed with moral inquiry. For any particular definition of lying, we can ask whether lying so understood is the subject of a perfect duty.

Moreover, notice that we have good reason to be skeptical about the prospects of settling on a single definition of lying as most intuitive. The data that drive these discussions often concern cases about which people have conflicting intuitions. And the foregoing discussion suggests a further complicating element in our intuitive reaction to definitions of lying. Speech is governed by a complex interplay of moral and linguistic norms. As we have seen, these norms may not always line up. But our intuitions about what counts as a lie are likely influenced by both sets of norms. So, our intuitions about

---

<sup>32</sup> Many have drawn a closer connection between these two kinds of norms, including Kant in at least some places. In the *Metaphysics of Morals*, Kant argues that we have a perfect duty to ourselves not to lie. And he writes, "communication of one's thought to someone through words that yet (intentionally) contain the contrary of what the speaker thinks on the subject is an end that is directly opposed to the natural purposiveness of the speaker's capacity to communicate his thoughts, and is thus a renunciation by the speaker of his personality, and such a speaker is a mere deceptive appearance of a human being, not a human being himself." Immanuel Kant, *The Metaphysics of Morals in Practical Philosophy*, Mary J. Gregor, ed. and trans. (New York: Cambridge University Press, 1996) 553. The account I describe in the text rejects this tight connection between linguistic and moral norms in order to leave room for permissible lies.

<sup>33</sup> See, for example, Carson (2006).

<sup>34</sup> See, for example, Carson (2006).

what counts as a lie may not be entirely coherent. For all of these reasons, it seems both possible and appropriate to conduct our moral inquiry without settling the definition of lying.

Notice, though, that the definitional debate is often conducted in conjunction with a debate about how to think about speech, e.g. how we ought to understand assertions. And, as I have indicated, settling these questions is incredibly important for moral inquiry because the moral norms, whatever they are, depend in a crucial way on the linguistic norms. So, in setting aside the definitional debate, I am not dismissing the significance of these other debates for moral inquiry.

### **3. Conclusion**

I have argued that the duty not to deceive is best understood as identical with or a part of the duty of doxastic concern. This is an imperfect duty that can be understood using the imperfect duty of beneficence as a model. We are required to take the permissible ends of others as reason giving. Similarly, we are required to take the agential interest in the truth that others have as reason giving. We have reasons to help others come to true beliefs and we have even stronger reasons not to interfere with their possession of true beliefs. As the extent of our interference increases, it becomes progressively more difficult to justify. This means that deceiving others through our behavior or our speech will often be impermissible.

Although the duty of doxastic concern suffices to condemn many lies, I have argued for the possibility of a distinct duty not to lie understood on the model of the perfect duty to keep one's promises. The analogy with the duty to keep one's promises enables us to see how there could be a distinct duty not to lie even if lying is not always morally worse than other forms of deception. Moreover, the example of the duty to keep one's promises enables us to see how a perfect duty can play a distinctive role in our moral reasoning without being implausibly strong. So, the result of this investigation is a framework for vindicating the commonsense view that there is something distinctively problematic about lying. A full defense of this view requires settling difficult questions about the linguistic norms that govern speech, but it does not require settling on a definition of lying.