10

Deception and the Nature of Truth

Michael P. Lynch

Philosophers who worry about the nature of truth rarely worry about the nature of deception. This is partly due to a sensible division of labor. When philosophers ask "What is truth?" they take themselves to be asking about a particular property of our beliefs. Defining that property—or as is fashionable, claiming that it can't be defined—is a matter of metaphysics and the philosophy of language. To talk about deception or lying on the other hand is to talk about the messy realm of human interaction. And that is the province of the ethicist. The general assumption seems to be that the philosophy of deception is to the philosophy of truth as bioethics is to the philosophy of biology. Nice stuff to know, undoubtedly important, but really not quite to the point.

To the layperson, this is apt to seem completely backward. Most folks, if they think about truth at all, think of deception first and truth second. Deception, after all, is a real human universal; it knows no boundaries. Most people are interested in what philosophers say about truth because they are fascinated by deception, not the other way around.

In this essay, I will try to say something about both deception and truth, with an eye toward vindicating the layperson's sense of what is important. I think that our attitudes toward deception tell us something important about both the nature and value of truth. These reflections in turn underline a more general lesson: that truth is a concept best understood in terms of the role it plays in our overall cognitive life.

What follows is organized as follows. In section 1, I lay out what I take deception to be, and argue that, unlike lying, it is conceptually linked to truth. In part two, I give some reasons for thinking that certain theories of truth are implausible in virtue of what they imply about deception. Since we arguably have a better grip on deception than on truth, this is bad news for such theories. In part four, I reflect on what our understanding of deception tells us about how and why we value truth. I concluded by drawing out a general lesson from these reflections about the possibility of giving a meaningful theory of truth.

10-Martin-Chap 10.indd 188 2/10/2009 1:37:52 AM

OUP • UNCORRECTED PROOF

1. DECEPTION

First there was the Word. And then there was the Lie. Lying is a public act, an act of speech. Deception on the other hand runs the behavioral gamut. As everyone knows, you can deceive without lying, indeed, without even speaking: sleight of hand, emotional misdirection, or simply pointing in the wrong direction ("He went that way") will do just as well. This line of thought encourages us to think of lying as a form of deception.

This is understandable, but mistaken. It is understandable since one can't define lying without appealing to deception. Lying isn't just saying what one doesn't' believe. Actors do that in the course of their profession, and they aren't lying, they are acting. To lie is to say what one doesn't believe with the intention of deceiving.[1] Lying and deception are conceptually linked in this way, but in fact lying is not a form of deception. For one can also lie but not deceive. Lying requires an intention to deceive, but what is intended may not happen: the listener may not *be* deceived. This might happen in two ways. First you might be skeptical and not believe what the liar says. Second, you might believe what he says but he inadvertently says what is true. In either case you were lied to, but you were not deceived.

It is this second case that is most important. Supposed I believe that the butler killed Col. Mustard in the library with the candlestick. Being a friend of the butler, I tell you what I think is false—that the butler was sunning himself in Hawaii at the time. You believe what I say, and take me to be sincere. But in fact I am mistaken, and the butler *was* in Hawaii, as you already know. I have lied: I've told you what I thought was false in order to mislead you. But I didn't succeed. I wanted to deceive you about the butler's whereabouts at the time of the foul deed in question, but you were not deceived about his whereabouts.

Of course, this is consistent with my deceiving you about something else. Since you believed that I was sincere, then I have deceived you about the fact that I wanted to deceive you. I've concealed my intentions, as we say. Indeed, any time I lie to you, and you don't detect the lie, I have deceived you about my lying to you. But this is an independent matter from whether I deceive you about the subject matter of the lie.

So it is wrong to say that lying is a form or kind of deception. One can lie without deceiving and deceive without lying. We might say that a successful lie is a kind of deception. But even here we must be careful. If by "successful lie" we mean an act that succeeds at being a lie, then as just

1. For a defense of this view of lying, see Michael P. Lynch, *True to Life* (Cambridge, Mass.: MIT Press, 2004). Compare Donald Davidson, "Deception and Division," in *The Multiple Self*, ed. J. Elster (Cambridge: Cambridge University Press, 1987), 79–93. A standard conception is Sissela Bok, *Lying: Moral Choice in Public and Private Life* (New York: Vintage, 1979); see also Bernard Williams's treatment in his *Truth and Truthfulness* (Princeton: Princeton University Press, 2002), 102–4.

10-Martin-Chap 10.indd 189 2/10/2009 1:37:52 AM

OUP • UNCORRECTED PROOF

indicated, lies, successful or not, don't always lead to deception. But if, as we probably should, we take "successful lie" to mean a lie that succeeds at doing what the liar intends, then successful lies are a type of deception. But not all lies, thankfully, are successful.

One way of putting the point I'm making here is that "deceive" is a success term. I deceive you only insofar as I actually succeed at misleading you or directing you away from the truth. Moreover, this misleading must be *willful* or *nonaccidental*. This is because those who give us wrong information by accident are not deceitful, but simply in error. In the admittedly technical way I will be using the term here, a willful action is one that is the result of a motivating reason, whether or not one is conscious of that reason. I can willfully mislead you without doing so with conscious intention: I may deceive you about Lancelot's devotion to you because I desire you for myself, even if I am unaware of my true motivations.[2] Likewise, I might deceive myself about your feelings by causing myself to turn away from the evidence simply because of my unconscious desire to want your feelings to be other than what they are.

With "willfulness" understood in this very broad way, we can suggest a particular constraint on deception:

> X deceives Y only if X willfully causes Y to have a false belief.

But this is not quite right. I can deceive you even if I don't cause you to believe something false.[3]

Consider the shell game. The con man presents three shells, one of which has a penny underneath. He moves the shells around and asks you to pick the shell with the penny. If done right, it looks easy, but isn't. The reason is that he distracts you (usually with subtle hand movements) so that you fail to track the right shell. This causes you to fail to know where the penny is. But one can lack knowledge without having a false belief. One can be simply confused, and that is typically the case with such tricks. You don't know what to think, and so simply guess. If so, then the con has succeeded—you've been deceived—because then the odds are in the con man's favor. This suggests that one can be deceived not only be believing what is false but by not believing what is true. That is:

> X deceives Y only if X willfully causes Y to fail to believe what is true.

But even this isn't quite right. Can't I deceive you into believing the truth? Suppose you believe falsely that Guinevere hates you when in fact she is madly in love with you. Pretending to be Guinevere, I write you a

2. For more on nonintentional accounts of deception, see A. Mele, *Self-deception Unmasked* (Princeton: Princeton University Press, 2002), and M. Johnston, "Self-deception and the Nature of Mind," in *Perspectives on Self-Deception*, ed. B. McLaughlin and A. O. Rorty (Berkeley: University of California Press, 1988); and J. Bermúdez, "Self-deception, Intentions, and Contradictory Beliefs," *Analysis* 60, 4 (2000), 309–19.

3. Thanks to Joel Kupperman for helping me to see this point.

OUP • UNCORRECTED PROOF

love note, you change your mind, and everyone lives happily ever after. I've deceived you but I've also caused you to believe what is true; but what I've deceived you about is not the truth I've caused you to believe. I've deceived you about who wrote the letter, but I've not deceived you about Guinevere's devotion. Nonetheless, this teaches us that deception is always deception with regard to something:

> X deceives Y with regard to f only if X willfully causes y to fail to believe what is true with regard to f.

For my purposes, we need only to treat this as stating a necessary condition. But I think there are good reasons to go further and treat the "only if" as a "if and only if." If we do, then we can say that to deceive is to prevent someone from grasping what is true about something. To put it differently, I deceive another when I *willfully cause them to be in a state of ignorance or error; I deceive myself when I do the same to myself*. And this in turn suggests a sway of saying what deception in general is: if we say for simplicity's sake that error is a form of ignorance, we can say that deception is *willful ignorance*.

2. TRUTH AND VULNERABILITY

That deception is willful ignorance tells us that when we are deceived, we lack a belief with a particular property—truth—about some matter before us. And that, in turn, suggests that our understanding of the nature of truth and deception are apt to be intertwined. A particularly stark way of illustrating this point is the simple consensus view of truth:

> S's belief that p is true if and only if S's belief that p is accepted within S's community.

This is the sort of position you end up with if you define truth in terms of warranted assertibility or belief, and then adopt Richard Rorty's position that "warrant as a sociological matter, to be ascertained by observing the reception of S's statement by her peers."[4] As Rorty once infamously put it, truth then becomes a matter of what your peers let you get away with.

One—among many—reasons to reject the simple consensus view is that it doesn't jibe with what we know about our ability to be deceived. If my earlier reflections on deception were correct, you are deceived only when someone willfully causes you to not believe what is true. This suggests that on the foregoing theory of truth, there are at least two ways you might deceive me: by causing me to lack a belief that is otherwise accepted in my community, or by getting the community to accept something I don't. Either way, you willingly cause me to not believe what is

4. "Putnam and the Relativistic Menace," *Journal of Philosophy* 90 (1993), 450.

10-Martin-Chap 10.indd 191 2/10/2009 1:37:52 AM

OUP • UNCORRECTED PROOF

true. But that seems off, to say the least. Suppose I don't believe you are honest but you wish to deceive me into thinking you actually are. If by use of a clever advertising campaign you convince everyone else in my community that you are honest, peer pressure may eventually cause me to doubt my previous belief; but if I am stubborn and don't change my mind, the mere fact *that you have convinced others won't make you honest*, or my assessment of your character mistaken. You can't deceive *me* by convincing *someone else* that something is the case.

If this weren't bad enough, the simple consensus theory entails that it is impossible to deceive the entire community. According to the view, if what is accepted within my community is p, then it is true that p. Thus, the community cannot be caused to be in ignorance or error about what they accept, for it is by definition true. But of course this is nonsense. Entire communities can and have been deceived about all sorts of things.

Most folks—we hope!—don't take the simple consensus view seriously. But the foregoing objections are worth noting because they illustrate how a theory of truth can go awry simply by way of what it implies about the possibilities for deception. Since we have a better grip on deception than on truth, a theory of truth that gets deception wrong is, to that degree, implausible.

The foregoing points also suggest another lesson: that there is a conceptual linkage between how objective an account takes truth to be and the extent to which the account makes us vulnerable to deception. The simple reason for this is that the twin hallmarks of objectivity are ignorance and error. What we believe may not be true, and what is true we may not believe. The more room an account leaves for ignorance and error, the more objective it intuitively counts as being. And since deception is the willful causing of ignorance and error, this suggests that the more vulnerable a theory of truth leaves us to deception, the more objective it is. But the linkage also goes the other way as well: the more objective a theory of truth, the more vulnerable it leaves us to deception.

To see this second point, consider radical deception. We are vulnerable to being radically deceived about some matter just when we are vulnerable to being caused to be in perpetual and undiscoverable ignorance or error about it. Consider a traditional correspondence theory according to which truth consists in correspondence with mind-independent fact. Such theories are nonepistemic: whether a belief corresponds to the facts does not depend in any way on whether we believe, justifiably or not, that it does. Humans do seem vulnerable to radical deception if truth is radically nonepistemic. For if truth has nothing to do with the epistemic status of our beliefs, then we might be deceived in trusting even our best theories of the world. We may be deceived by the Cartesian demon, or be brains in vats. Hence, it may seem that we can't be sure that we really know what we think we know. Thus the usual complaint against correspondence, nonepistemic theories of truth: they make us vulnerable to

10-Martin-Chap 10.indd 192 2/10/2009 1:37:52 AM

skepticism because they make us vulnerable to the possibility of radical deception.[5]

Antirealist theories of truth have been traditionally motivated by this complaint. Thus, suppose we hold a Peircian view like:

(P) <p> is true if and only if <p> would be accepted at the end of inquiry.

Radical deception of the Cartesian variety is ruled out by fiat by such a view. For while we might be deceived about what will or will not be accepted at the end of inquiry, what we do in fact accept at the end of inquiry must be true. Even the evil demon cannot deceive us at the hypothetical limit of science, because truth is defined as what we believe when we reach that limit. What we believe at the ideal limit can't be wrong. This means that there is less space for ignorance and error on this account. Intuitively, therefore, we might say that the Peircian view is more objective than the simple consensus view, but less objective than the correspondence view canvassed earlier. And one reason for this is that it allows greater room for the possibility of deception, although it rules out radical deception. And this in turn suggests the other direction of the aforementioned conceptual linkage. Not only is it the case that the more vulnerable a theory of truth leaves us to deception, the more objective it is, the less vulnerable to deception an account leaves us, the less objective we will take it to be.

The point holds even for antirealist views of truth that are not motivated just by fear of skepticism. Consider, for example, a theory of truth built out of a notion like superwarrant:[6]

Superwarrant: <p> is superwarranted if and only if the belief that p is warranted at some stage of inquiry and would remain warranted at every successive stage of inquiry.

Here a "stage of inquiry," as the name suggests, is a state of warranted information or evidence available in principle in the actual world to some open-minded, receptive inquirer. Stages are understood as being extensible (additional information might always come in) and inclusive (the additional information is just that—additional; all successive stages of inquiry include the information warranted at prior stages). Again, superwarrant does not posit an idealized "end of inquiry." A superwarranted belief is one that is warranted by some state of information available to *ordinary inquirers* and that, in fact, would never be defeated or undermined by subsequent increases of information also available to ordinary inquirers.

5. For a similar point lining skepticism and realism, see J. Heil, "Mind and Knowledge," in *Oxford Handbook on Epistemology,* ed. P. Moser (Oxford: Oxford University Press, 2002), 316–35.

6. See Wright's *Truth and Objectivity* (Cambridge, Mass.: Harvard University Press, 1992). His term is "superassertibility." I use "superwarrant" both because my account of the notion is a bit different from his own, and because I find it a more perspicuous label.

OUP • UNCORRECTED PROOF

Moreover, superwarrant is a stable property: if a belief is warranted, then it is superwarranted at any stage of inquiry. Thus we might suggest:

(SW) A proposition is true if and only if it is superwarranted.

This claim is clearly distinct from (P). Nonetheless, it has some of the same consequences. In particular, it rules out the possibility of radical deception. For (SW) and our account of superwarrant seemingly underwrite the claim that

If <p> is superwarranted, <p> is knowable.

And consequently, we can deduce that

If <p> is true, then <p> is knowable.

Now given the T-schema, or the principle that

<p> is true if and only if p,

we can infer that

If p, then <p> is knowable.

This is just to draw out the obvious: that any account of truth like (SW), which defines truth partly in terms of warrant or a belief's epistemic status, will have to admit that truth is "epistemically constrained." But now it is a simple matter to show that any such theory limits the possibilities for radical deception. For if we accept that if p, it is knowable that p, then presumably we should also accept the following:

If <p> is knowable, then p is not something about which reflective humans may remain perpetually and undetectably ignorant.

But if the fact that p is not something about which reflective humans may remain forever ignorant, then it is not something about which we can be eternally deceived. We are immune from radical deception because at some point, to someone, the truth will out, so to speak. In short, (SW) appears to imply that

If p, then it is impossible to be radically deceived about p.

This point exposes the weakness of (SW) as a theory of truth. For it seems very likely that there *are* some truths about which we cannot rule out, a priori, the possibility that we might be radically deceived about them. We don't need evil demons to make this point either. It seems possible that we might be perpetually and undetectably deceived about some event in the distant past by certain documents that were willfully created for that purpose. We will never have any means by which to see through the deception, or even any evidence that it is a deception. If this is possible, then there are at least some propositions about which it is possible to be radically deceived.

This suggests that superwarrant is not a plausible theory of the nature of truth. That is, it is not plausible that truth just is superwarrant. But if

10-Martin-Chap 10.indd 194 2/10/2009 1:37:52 AM

OUP • UNCORRECTED PROOF

we were to stop here, we would overlook and important point. For even if truth can't be identified with superwarrant, it may still be plausible that some propositions are *made true* by being superwarranted. Indeed, this is just the sort of possibility that a pluralist theory of truth allows for.[7] Such theories, at a minimum, are committed to the idea that

> **(ST):** Necessarily, for any proposition, if it is true, then it has some property F such that, necessarily, if a proposition if F, it is true.

What (ST) proposes is that truth is a *supervenient* property, in that it strongly covaries with other properties. And (ST) is compatible, clearly, with two further thoughts: first, that truth does not just covary with these further properties but is metaphysically dependent in some way on them, and second, that which property determines truth can *vary*. That is, not only:

> Necessarily, there are some propositions such that if they are F, then they are T

but also:

> It is possible that there are some propositions that are T but not F.

This limited pluralism about the base properties for truth says nothing about the property of truth itself *other* than it is a single higher-level property that is *asymmetrically dependent* on other properties. Intuitively, these other properties are those that *make propositions true*. Perhaps for some sorts of propositions, the property that makes them true is superwarrant. Indeed, the foregoing reflections help to show us which sorts of propositions are likely to be candidates: propositions about which it is a priori the case that we cannot be radically deceived. What sort of propositions might these be? One suggestion—and here I only put this forward as a representative suggestion—would be moral propositions. Many writers, holding quite different normative ethical theories, think moral wrongness is conceptually tied to responsibility and blameworthiness. "We do not do call anything wrong," as Mill writes, "unless we mean to imply that a person ought to be punished in some way or other for doing; if not by law, by the opinion of his fellow creatures; if not by opinion, then by the reproaches of his own conscience."[8] What is wrong is what we can be justifiably held responsible for doing—what is worthy of blame, in short. If so, then the following argument seems cogent:

7. For two versions of alethic pluralism see Wright, *Truth and Objectivity* and *Saving the Differences* (Cambridge, Mass.: Harvard University Press, 2004); Michael Patrick Lynch, *Truth as One and Many: A Pluralist Manifesto* (Oxford: Oxford University Press, 2009); and "Truth and Multiple Realizability," *Australasian Journal of Philosophy* 82 (2004), 384–408.

8. Mill, *Utilitarianism*, ch. 5. See also Stephen Darwall, *The Second-person Standpoint*, (Cambridge, Mass: Harvard University Press, 2006), 27. Robert Adams, *Finite and Infinite Goods* (New York: Oxford University Press, 2002), 233–38; Allan Gibbard, *Wise Choices, Apt Feelings* (Oxford: Oxford University Press, 1990), 42; John Skorupski, *Ethical Explorations* (Oxford: Oxford University Press, 2002), 142.

10-Martin-Chap 10.indd 195 2/10/2009 1:37:53 AM

OUP • UNCORRECTED PROOF

If an action is wrong, then it is blameworthy.

If an action is blameworthy then it is knowable that it is wrong.

Therefore, if an action is wrong, then it is knowable that it is wrong.

The argument appears to be valid. The first premise is Mill's conceptual truth. The second premise, too, seems to be a conceptual truth about blameworthiness. It says that it is a necessary condition of being a blameworthy action that it is possible that some ordinary human observer could at some point recognize it as wrong. My action is blameworthy when it is possible for me to be held accountable for that action. But I can't be held accountable, surely, for an action that *no one, including myself*, would ever know that it is one for which I should be censured. Given the first premise, to think otherwise would simply mean not only that there are unnoticed moral wrongs but also that there are unnoticeable moral wrongs. It might, for example, be wrong that I used more that six words in this sentence. And that seems absurd. Hence, the conclusion seems warranted, and so, presumably, would be a parallel conclusion about praiseworthiness and rightness. If so, then, thanks to the schema that it is true that p if and only if p, we know that there are no unknowable truths about what is morally right or wrong.

Similar reasoning shows that radical deception about moral wrongness is deeply implausible. Such deception would be possible if it were possible that an evil demon could make us believe that it is permissible that I write this sentence when in fact it is morally wrong for me to do so. But if no one is ever, even if in principle, able to detect the demon's deceit, then on one could ever know that my action in writing that sentence was morally wrong. But if so, that is, if no one is ever able to know it is wrong, then by the foregoing argument, it isn't wrong.

Of course, the fact that I can't be radically deceived about the moral truth doesn't mean that I can't be substantially deceived. The fact that something is unknowable in principle is consistent with its being the case that I and everyone on the face of the earth are deceived about it right now. Radical deceit may be ruled out for morality, but global deceit is not. Nonetheless, this fact about moral truths makes them noticeably different from truths about the natural world, where we are quite willing to accept unknowable truth and radical deception to boot.[9]

These reflections help to illustrate not only the conceptual linkage between deception and the nature of truth, but the general lesson I drew

9. Obviously the foregoing considerations aren't, all by themselves, intended to *prove* that moral judgments are made true by being superwarranted, or even that moral truth is epistemically constrained. Extreme utilitarians might resist the argument by appealing to the possibility of incalculable utility functions. Others might acknowledge its force but insist that it is a sure fact about the world and our cognitive capacities that we happen to be able to recognize the moral facts when they obtain. Lucky devils that we are, we are just built to be able to discover such facts. For my part, I see no reason to think the world is so cooperative. I take it to be more plausible that the foregoing argument tells us something about how those beliefs are made true.

OUP • UNCORRECTED PROOF

out of that linkage earlier. Where truth is maximally objective, we are vulnerable to radical deception. It is plausible that we are vulnerable to such deception, and therefore theories of truth that artificially limit the nearly limitless ways humans might be in ignorance or error, hence the ways they may be deceived, should be rejected. But where the possibility for radical deception is in fact already limited, such theories may in fact be plausible accounts of what makes some propositions true.

3. DECEPTION AND THE VALUE OF TRUTH

In the previous section, I reviewed some of the ways the conceptual linkage between deception and truth can, and should, affect how we think about the metaphysics of truth. I now turn to question of how that linkage affects how we think about the value of truth. Since I've addressed this issue at length elsewhere, I'll be brief.[10]

People don't like to be deceived—by their mechanic, their boss, their friend, their lover, anyone. One reason for our dislike is obvious: when you are deceived, someone has caused you to not believe what is true, or even to believe what is false. Either way, they've willfully placed themselves between you and the facts. And that can be a dangerous thing. But even if it is not overtly dangerous, it is likely to interfere with one's plans, to change them, or to simply make them go awry. Either way, we are less likely to get what we want. And that of course is typically the point: the deceiver deceives to get his way, not (generally speaking) to facilitate yours. There are exceptions of course, but generally speaking, we hate to be deceived because deception has negative practical consequences.

But our dislike of deception goes deeper than that. One reason to think so is that most folks would prefer not to be deceived even if it would make no difference to one's experiences. Suppose you had to choose between two doors. Once you make the choice of which door to enter, you will forget that you ever made a choice at all. Behind door number one is your life just as it is now. Your friends are friendly, and your lover loves you. Behind door number two is a very similar life, with one very important exception: here some of your friends and your lover really despise you. But you will never discover that fact: their deceit will be perfect. From the inside, both lives will be indistinguishable: where the first causes you joy, the second does also; where the first causes you pain, the second does as well; and so on to the grave. Yet in the second, your life is the life of the fool: you are deceived.

Forced to choose, almost all of us will prefer the first life over the second. Perhaps some may be ambivalent; they'll flip a coin. Presumably, no one will actively prefer the second over the first. Either way, your reaction tells you something about how deeply you dislike deception. If you

10. See *True to Life*, and "Replies to Critics," *Philosophical Books* 46 (2005), 331–42.

10-Martin-Chap 10.indd 197 2/10/2009 1:37:53 AM

OUP • UNCORRECTED PROOF

are ambivalent, then deception matters less to you than it does to others. What matters is how you feel and experience life; if the truth of your beliefs has no effect on those feelings and experiences, then you don't care whether you are being deceived or not. One door is as good as another. Most of us, however, will find this attitude odd, even repugnant. We don't just want to seem to have friends and lovers, we want friends and lovers, even if there were no discernible difference between the one case and the other. Moreover, we want to want to be that way: we care about not being deceived. We would no more wish to be ambivalent about which door to enter than we would wish to willingly enter into a deception.

Our attitude toward such choices also tells us something about our attitudes toward truth. Given the conceptual connection between truth and deception, this is not surprising. The fact that we prefer not to be deceived—even when the deception is undetectable—suggests that our preference for believing whatever is true over not doing so remains even when it would have no effect on how we experience life. And this in turn suggests that it is a basic preference—not derived from a preference for something else. Moreover, it is not a mere preference—I don't just want to believe whatever is true; I care about doing so. And if I care about not being deceived for more than instrumental reasons, I care about believing what is true for more than instrumental reasons. I care about it "for its own sake."[11]

It is worth emphasizing that this line of reasoning is not intended to show, absurdly, that we want all of our actual beliefs to be true. I believe many propositions that I don't want to be true. Beliefs about the future of global warming or the continuing spread of AIDS in Africa number among them. But the fact that I don't want these particular propositions to be true is entirely consistent with it being the case that I care about believing what is true and only what is true, whatever that turns out to be.

Nor does our disvaluing of deception and consequent valuing of true belief mean that truth is our only or ultimate value. It obviously is not. Sometimes other things matter more than truth. Thus, more of us would be willing to be deceived, or to deceive ourselves, if we thought that more good than bad would come of it overall, or that the matter was so trivial that the point was essentially moot. But this fact is entirely consistent with the fact that considered by itself, deception is still something we deeply wish to avoid, and believing what is true is something we care about achieving. What this shows is that, like almost everything else we care about, true belief is a pro tanto value. It is something we care about other things being equal.

11. What we care about is the state of affairs of believing what is true. This value or end is to be distinguished from the fact that it is correct to believe a true proposition—or that the standard of correctness for belief is truth. See Lynch, *True to Life*, and "The Truth of Values and the Values of Truth," in *Epistemic Value,* ed. D. Pritchard (Oxford: Oxford University Press, forthcoming).

OUP • UNCORRECTED PROOF

What we care about we typically believe is worth caring about. Thus, what these reflections do show is that we (or most of us) believe that truth is worth caring about, and caring about for its own sake.[12] And that in turn tells us that truth is a value: the state of affairs of believing what is true is something we take to be a good.

4. CONCLUSION

Philosophers like me who think about truth for a living are, funnily enough, often the least likely to believe that there is anything all that interesting to say about truth. One reason for this is that most of us know too well the failures of traditional theories like the correspondence theory or Peirce's pragmatic view. Such theories are reductive: they try to identify truth with some single property that all and only true propositions have in common. But such attempts are beset by counterexamples. Theories that seem plausible when applied to propositions about the middle-sized dry goods of everyday life (like certain versions of the correspondence theory) seem much less plausible when applied to propositions about abstract entities like numbers or norms. And theories that seem more plausible when applied to propositions about norms (such as superwarrant) are much less plausible when applied to propositions about the physical world around us.[13] Partly as a result, many philosophers working on truth today are attracted to one form or other of deflationism, according to which, roughly, everything that needs to be said about truth can be gleaned from our inclination to accept instances of the T-schema.[14] From this standpoint, it is simply a mistake to think with the traditional theories that truth has any sort of "nature" that is worth explaining. It is a useful concept for sure, providing a handy semantic ladder by which we can ascend and generalize over infinite strings of propositions, but it is not to be confused with a property that needs deep metaphysical investigation.

In my view, the considerations raised in this article suggest that there may be more to say about truth than jaundiced deflationists believe, even if it isn't the sort of thing traditionalists look for. What the foregoing thoughts suggest is that truth is a concept that is intimately related to a host of other concepts—deception, ignorance, objectivity, value. Truth seems in fact to sit in a network of such interrelations. Consequently,

12. Of course, this is distinct from showing that truth actually is worth caring about. For arguments to that effect, see ch. 8 of *True to Life*.

13. For further arguments to this effect, see < Lynch, "Truth and Multiple Realizability."

14. Representative deflationists include P. Horwich, *Truth*, 2nd ed. (Oxford: Oxford University Press, 1998), and H. Field, *Truth and Absence of Fact* (Oxford: Oxford University Press, 2001). For some remarks that overlap with the foregoing, see Donald Davidson, "The Folly of Defining Truth," in *The Nature of Truth*, ed. M. P. Lynch (Cambridge, Mass.: MIT Press, 2001).

10-Martin-Chap 10.indd 199 2/10/2009 1:37:53 AM

OUP • UNCORRECTED PROOF

there is reason to think that revealing these interrelations will shed light on the concept of truth, by illuminating what can be described as the functional role truth plays within our overall cognitive economy. Investigating truth's role will most likely not aid us in a reductive analysis of truth. It will not reveal the secret essence of truth. But it will tell us more about what truth does for us—how it functions in our thought—not just in logic or in epistemology, but in the broader realm of messy human interactions, the realm that the layperson lives in, the realm we live in ourselves.

ACKNOWLEDGMENTS

Thanks to the following individuals for helpful discussion: Paul Bloomfield, David Capps, Daniel Massey, Joel Kupperman, and Terry Berthelot.

10-Martin-Chap 10.indd 200 2/10/2009 1:37:53 AM