Lying, risk and accuracy

SAM FOX KRAUSS

A large literature has yielded near unanimity on two necessary conditions on lying.[1] One lies about p only if one:

(1) Says that p.
(2) Believes that not-p.

Philosophers have discussed what satisfies the first necessary condition and what, if any, further necessary conditions there are. But, absent from the literature is any serious discussion of the second necessary condition – what I'll call the *belief requirement*. In addition, although epistemologists have long recognized that people have attitudes of *degrees* of belief instead of (or in addition to) full belief, philosophers have not considered the possibility that the true requirement on lying concerns not the speaker's belief, but rather her degree-of-belief.[2]

I begin with a claim about what makes liars blameworthy, when they are. I claim that the correct explanation of why liars are blameworthy includes the liar's imposing a risk on the audience. This blameworthiness admits of degrees: the greater the liar's confidence that what she asserts is false, the greater the risk she'll think she's imposing on the dupe, and, therefore, the greater her blameworthiness. From this explanation of what makes liars blameworthy, I arrive at a dilemma: either the belief requirement is wrong,

---

1 See, e.g. Chisholm and Feehan 1977: 152; Meibauer 2005: 1376; Fallis 2009: 33; Saul 2012; Lackey 2013: 246; Stokke 2013: 33; Webber 2013: 652; Rees 2014: 59; Leland 2015: 550; Keiser 2016: 462; Michaelson 2016: 481.

2 I'm not the first, however, to consider the role of degrees of belief in the context of lying. In response to a claim by Sorensen (2010) that knowledge-lies do not involve deception, Staffel (2011) and Fallis (ms) argue that a dupe is deceived by a liar when the dupe's confidence in a falsehood increases, even if the dupe does not thereby come to *believe* the falsehood. Chisholm and Feehan (1977) mention that the extent of deception depends on the degree of belief the dupe comes to have in the falsehood the liar asserts. While these authors offer important precedents for including degree-theoretic reasoning in accounts of lying by thinking about the extent to which the dupe is deceived, my project is to discuss the *speaker* using degrees of belief, and, importantly, to give an account of blameworthiness for lying that depends on the relationship between the speaker's degree of belief and her expectation about what the dupe's degree of belief will be. Thanks to an anonymous referee for urging me to clarify this point.

   Marsili (2014) is, as far as I can tell, the sole person to argue for an account of lying that explicitly takes into consideration the speaker's degree of belief. He replaces the belief requirement with a requirement that the speaker must be more confident that the proposition she asserts is false, than true. In Section 3, I show why this suggestion is not correct.

or *lying* isn't a particularly interesting concept. I then suggest a principled replacement to the belief requirement: the *worse-off requirement*.

1. The simple reason liars are usually blameworthy

Philosophers who theorize in terms of full beliefs may think about the blameworthiness of liars in the following way. A rational agent uses her beliefs, in combination with her desires, to act in order to satisfy her desires (see, e.g. Ramsey 1926/1990: 68–74; Stalnaker 1984: 16–18; Dennett 1987: 20–33, 43–57). When a person has true beliefs, she acts in ways that tend to satisfy her desires; when a person has false beliefs, she acts in ways that tend to frustrate her desires. Thus, we care about having true beliefs. (This is consistent with acknowledging that it may not be the case that it's *always* better for us, practically, to have true beliefs, see, e.g. Kelly 2003: 624–5.) By lying, the speaker expects to cause the dupe[3] to have a particular false belief: the proposition she asserts. When the dupe uses that belief in her practical reasoning, she won't get what she wants. I agree that this is an appealing explanation of why lying is wrong.

A degree-theoretic account of belief, however, allows for greater precision and a finer explanation. We can say, instead, that a person has a *credence* in a proposition, where a credence is a subjective probability one assigns to the truth of a proposition and ranges from 0 (the minimal degree of confidence) to 1 (the maximal degree of confidence). While beliefs are either true or false, credences are more or less *accurate*. Accuracy is measured by a scoring rule, which takes as its inputs the agent's credence and the truth of the proposition. We might think of the accuracy of a credence as a measure of 'its "closeness to the truth"' (Schoenfield forthcoming: 5). For example, certainty of a true proposition is maximally accurate and certainty of a false proposition is maximally inaccurate. Being unsure of a true proposition is more inaccurate than if one were sure, but less inaccurate than if one were sure of its negation. I'll note that there is significant controversy about the appropriate desiderata for epistemic scoring rules.[4] For purposes of this article, though, we can abstract away from most issues that cause controversy. We need only assume what is widely agreed upon: that accuracy scores of credences increase as the credence approaches the truth-value of the proposition.

Because we need accurate credences to act in ways that tend to satisfy our desires, we care about having accurate credences. By lying, the speaker does

---

3   As best I could determine, the first use of *dupe*, in the philosophical literature, is due to Bok (1978).

4   See the literature on epistemic scoring rules, e.g. Joyce 1998; Schoenfield 2015; Konek and Levinstein forthcoming and references therein.

something that she expects, if she is trusted, will cause the dupe to have a more inaccurate credence in the proposition she asserts. This, I argue, is the better explanation for why people are blameworthy for lying.

Call the increase in inaccuracy in p the dupe suffers by trusting the liar about p *epistemic damage* and the increase in inaccuracy the liar expects the dupe to suffer *expected epistemic damage*.

2. Epistemic damage and expected epistemic damage

Consider an instance of lying from the perspective of an omniscient third party: the speaker has some small credence in a false proposition that she nevertheless asserts to the dupe. If the dupe trusts the speaker, the dupe will count the speaker's assertion as evidence for the proposition asserted, and the dupe's credence in that proposition will increase. Because the dupe has become more confident in a false proposition, she is epistemically worse off – she has suffered epistemic damage. The greater the inaccuracy of our credences, the less we will tend to get what we want. So, epistemic damage correlates with real-world, or, *practical* harm. (This is consistent with acknowledging both that epistemic damage does not always cause practical harm and that lying does not always lead to epistemic damage.)

Now consider an instance of lying from the liar's perspective: the liar has some low credence in a proposition that she nevertheless asserts to the dupe. The liar expects the dupe to trust her, and therefore expects the dupe's credence in the proposition to increase. From the liar's perspective, her own credence maximizes expected accuracy. Therefore, according to the liar, anyone whose credence moves farther away from her own has become more inaccurate. While there is controversy about this, the general consensus in the literature is that you ought to be *immodest* about the accuracy of your own credences. This means that you should think that the expected accuracy of your own credences is greater than the expected accuracy of any other credences you might adopt.[5] The intuitive thought is that it is irrational to (i) think that there are different credences with a higher expected accuracy and (ii) not adopt them as your own. Assuming immodesty, I give an account of lying in terms of epistemic damage, which is superior to existing accounts.

Note that while epistemic damage depends on the truth-value of the proposition the liar asserts, *expected* epistemic damage does not. For example, if the liar mistakenly asserts a true proposition believing it false, and the dupe increases her credence in the proposition, the dupe does not suffer epistemic damage because she is epistemically better off. Nevertheless, there is *expected* epistemic damage. In addition, the expected epistemic damage will depend on the extent to which the speaker thinks the liar trusts her. If, for example, the

---

5 See the literature on epistemic immodesty, e.g. Horowitz 2014 and references therein.

speaker knows that the dupe doesn't trust her at all, then there won't be expected epistemic damage. Note also that the expected epistemic damage is not the same as the liar's prediction about the dupe's change in credence. Because the liar's statement does not vary according to her confidence, we should expect the dupe's change in credence to be independent of the liar's credence. However, when the liar is *more* confident that what she asserts is false, she'll expect the dupe to become *more* epistemically worse off than in the case in which she is not as confident that what she asserts is false. Expected epistemic damage is a measure of the liar's estimate of how badly off the speaker will be, which is relative to whatever the liar thinks is the (maximally) rational credence to hold: her own. Expected epistemic damage is, of course, to be calculated using the speaker's degree of belief.

I claim that the greater the expected epistemic damage, the greater the liar's blameworthiness. The more confident the liar is that what she asserts is false, the more inaccurate she will think the dupe will become. The more inaccurate the liar thinks the dupe will become, the more likely the liar will think the dupe will not get what she wants. Once we adopt a degree-theoretic model of belief, it becomes clear that expected epistemic damage will vary according to the confidence the liar has that what she asserts is false. If we accept that epistemic damage correlates with practical harm, then it becomes clear that expected epistemic damage is a prediction of how badly off the liar thinks that her lie will make the dupe. In this way, a liar's blameworthiness is explained by the risk she imposes on the dupe. Furthermore, this blameworthiness admits of degrees: the greater the liar's confidence that she asserts something false, the greater the risk she'll think she's imposing on the dupe, and, therefore, the greater her blameworthiness.

3. Dilemma

By lying, the speaker does something that she believes will cause the dupe to become less epistemically accurate and, therefore, less likely to be able to satisfy her desires. I've argued that once we think in terms of degrees of belief, we can appropriately explain the blameworthiness of lying as the causing of expected epistemic damage, which depends, among other things, on the confidence the speaker has that what she asserts is false. But, once we so ground the blameworthiness of lying in expected epistemic damage, the orthodox account of lying faces a dilemma: either the belief requirement is wrong or *lying* isn't a particularly interesting concept. In what follows, I explain why the dilemma exists and suggest why impaling ourselves on the first horn isn't all that bad.

I've argued that blameworthiness for lying is best explained by the liar's imposition of risk on the dupe, which I analyse in terms of expected epistemic damage. If the belief requirement is correct, then there are nearby cases of

non-lies in which the speaker is blameworthy for the same reason as the liar, albeit to a lesser degree, without lying. However, it's difficult to see why the belief requirement marks the distinction between lies and assertions that, by failing to meet the belief requirement, fall just short of lying. I'll propose, in the next section, that we can make the distinction more intuitive if we abandon the belief requirement and instead introduce a requirement involving degrees of belief.[6]

To get a sense of why the belief requirement is not natural or intuitive, imagine a case in which someone is not so confident that we would ascribe *belief* to her – say, she has 0.6 credence that it will rain. If the speaker nevertheless says that it won't rain, the speaker will think that her audience will be worse off, for trusting her. This is the case even if the speaker's confidence is not sufficient for belief, and therefore, according to the orthodox account, not sufficient for lying. But, we can explain what's wrong about this assertion in the same way as what's generally wrong about lying. The speaker's assertion that it won't rain when she has 0.6 credence that it will, and an assertion that it won't rain when she *believes* that it will differ only by the magnitude of risk imposition, and, therefore, the speaker's blameworthiness. If this is the case, it's hard to see why *lying* picks out any special phenomenon.

Imagine if we were to define a 'theft' as an intentional taking of someone else's property worth at least $100. One might balk at the $100 requirement and ask for a principled reason why a $40 taking shouldn't count as a theft, too. Perhaps there is a pragmatic explanation: perhaps a statute requires police to investigate all 'thefts' but that it's inefficient to investigate takings of less than $100. So, defining 'theft' in this way makes pragmatic sense. One might think that there is some similar, philosophically pragmatic reason for the belief requirement, but it's far from clear that there is a philosophical equivalent to the pragmatic classification of legal concepts, and, even if there were, what that justification would be in this case. If the belief requirement is correct, then lying does not pick out an especially interesting concept.

To address the first horn of the dilemma, if lying really does pick out an interesting concept, then the belief requirement must be wrong. Marsili (2014) proposes we replace the belief requirement with the 'comparative insincerity condition', according to which the speaker must be more confident that the proposition she asserts is false than true. However, this doesn't

---

6 Others have offered different distinctions between lies and blameworthy assertions that fall short of lying. Benton (forthcoming), for example, argues that we ought to replace the belief requirement with a knowledge requirement. On his view, even if the knowledge requirement is not met, and thus one doesn't lie, one can be blameworthy if one 'negligently' asserts a proposition, where one negligently asserts a proposition just in case 'one doesn't consider whether one knows p' (11).

account for the damage lying does. Imagine a case in which the speaker is just slightly more confident that p is true than false – say she has 0.51 credence in p. And, further, imagine that the speaker knows that her audience is agnostic – that the audience has 0.5 credence in the proposition. From the speaker's perspective, the audience's credence is very close to her own, and, therefore, the speaker thinks the audience's credence nearly maximizes expected epistemic accuracy. When the speaker asserts that p, supposing the speaker believes the audience will trust her, she'll think that the audience will assume that she is far more confident than she is. Imagine that the audience moves from 0.5 credence to 0.8 credence. If the speaker's credence is 0.51, then the speaker will expect the audience to suffer epistemic damage. If the blameworthiness for lying is grounded in expected epistemic damage, then this speaker is blameworthy in exactly the same way as liars are, even if, according to both the orthodox position and Marsili's proposal, she hasn't lied.

4. The worse-off requirement

With the belief requirement, lies are not interestingly distinct from nearby cases. My proposal would save lying from being uninteresting. Using the simple, general explanation I've given for why liars are blameworthy, it's possible to construct a principled replacement to the belief requirement:

> *The worse-off requirement*
> The expected epistemic damage to the audience, with respect to p, by the speaker's lights, conditional on the audience trusting her with respect to p, at all, is greater than 0.

That is, it's a necessary condition on lying that the speaker think that, if the dupe trusts her, the dupe will be worse off, epistemically. I've deliberately made the requirement conditional on the audience trusting the speaker; it puts off a decision about whether intent to deceive is a necessary condition on lying. For those who argue that bald-faced lies (assertions by a speaker who believes both that the proposition she asserts is false and that her audience won't trust her) *are* lies, then the worse-off requirement and that a speaker says that p might be sufficient for lying.[7] Those who argue that intent to deceive is a necessary condition on lying would add a further condition: that the speaker expect the dupe to trust her. My worse-off requirement takes no position in this debate.[8]

---

7 See the literature on bald-faced lies, e.g. Sorensen 2007, Lackey 2013, Fallis 2015 and references therein.

8 This replacement to the belief requirement is consistent with the classification of some assertions as lies that the traditional account would exclude. In addition, some assertions that might have been classified as lies on the traditional account, no longer would.

My view allows that some assertions that are not lies can be blameworthy (of course), and I also allow that the blameworthiness of some of those assertions can be given a closely related, structurally similar explanation using the notion of expected epistemic damage. Where I draw a principled, natural boundary around lying, however, is when damage is done with respect to p, the proposition asserted. I claim this is more intuitive than where advocates of the belief requirement must draw the boundary.

I'll note here two further features of the worse-off requirement. First, it is compatible with epistemic models that eschew belief entirely in favor of degrees of belief.[9] If the belief requirement is correct, it's not clear what lying is for a Bayesian without belief, for example. My replacement to the belief requirement provides an answer. Second, the worse-off requirement is consistent with including knowledge-lies as lies. According to

---

An anonymous referee raises the following proposed counterexample to the worse-off requirement:

The speaker, A, has .6 credence in p, and knows that the audience, B, has .9 credence in p. A has good reason to believe that, because of the nature of B's evidence, B won't decrease her credence much below .6, even if A says that p is false. In a charitable effort to bring B's credence closer to A's own, and knowing that it won't fall much below .6, A says that p is false.

Assuming that .6 credence in p is sufficient for belief, most, I think, still wouldn't count this as a lie. That the liar must intend to deceive the audience is a view shared by most who've written on lying, though not all (see, e.g. Carson 2006; Sorensen 2007; Fallis 2009). In the case as described, A's intent is to make B's credence *more* accurate. And, A has good reason to think that, given the unusual and somewhat artificial nature of the thought experiment, by asserting that p, this will be the case. Thus, A does not intend to deceive B. So, assuming the plausible intent-to-deceive requirement, this is not an instance of lying. At worst, this case shows that the worse-off requirement works only on views according to which intent to deceive is a necessary condition on lying. However, I've intentionally made the worse-off requirement consistent with accounts of lying that do not require intent to deceive, so it'd be better for my purposes to not make this concession. So, let me add a further consideration.

The case as described ought to be counted as a lie only if we adhere to the belief requirement and don't require intent to deceive. The wide array of accounts of lying that include the belief requirement, I think, fail to properly appreciate the normative properties of lying. I've argued that a good account of lying ought to capture why liars are blameworthy. The belief requirement leaves out a great number of these kinds of cases. The worse-off requirement explains, better than the belief requirement does, what makes liars blameworthy – namely, that the liar expects her audience to be epistemically worse off. If this unusual case seems intuitively to count as a lie, then I think it's a justifiable cost to my proposal that it excludes this case in exchange for the great gain of many cases that share the interesting features of lies, but, on the traditional view, are excluded.

9   See, e.g. Jeffrey 1970, who writes that 'our ordinary notion of *belief* is only vestigially present in the notion of degree of belief' (171–72); or see Christensen 2004, who demotes the importance of binary belief for epistemic rationality (96–105).

Roy Sorensen,

> [a]n assertion that p is a knowledge-lie exactly if intended to prevent the addressee from knowing that p is untrue but is not intended to deceive the addressee into believing that p (Sorensen 2010: 610).

Sorensen claims that knowledge-lies do not involve deception, but Staffel (2011) and Fallis (ms) convincingly show that the interesting cases do, because even though the dupe does not come to *believe* the false assertion, she becomes more confident in it. The worse-off requirement is consistent with an account of lying according to which most, though perhaps not all, knowledge-lies count as lies.

I've argued that once we think in terms of degrees of belief and accept a simple explanation for the blameworthiness of lying, we arrive at a dilemma. To tackle that dilemma, I've suggested a principled alternative to the belief requirement – the worse-off requirement. Although nearly every philosopher who has given an account of lying has endorsed the belief requirement, my revision is not all that radical. The worse-off requirement recasts lying as involving an imposition of risk on the audience and better captures what usually makes lying blameworthy: that the speaker expects the dupe to be worse off, for trusting her.[10]

The University of Texas at Austin
Austin, TX 78712, USA
sfk@utexas.edu

References

Benton, M. Forthcoming. Lying, belief, and knowledge. In *The Oxford Handbook of Lying*, ed. J. Meibauer. Oxford: Oxford University Press.

Bok, S. 1978. *Lying: Moral Choice in Public and Private Life*. New York: Random House.

Carson, T. 2006. The definition of lying. *Noûs* 40: 284–306.

Chisholm, R. and T. Feehan. 1977. The intent to deceive. *Journal of Philosophy* 74: 143–59.

Christensen, D. 2004. *Putting Logic in Its Place*. Oxford: Oxford University Press.

Dennett, D. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.

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

Fallis, D. 2015. Are bald-faced lies deceptive after all? *Ratio* 28: 81–96.

Fallis, D. (ms). Bayesians don't tell knowledge-lies (and probably nobody else does either). Unpublished manuscript. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1666545

Horowitz, S. 2014. Immoderately rational. *Philosophical Studies* 167: 41–56.

Jeffrey, R. 1970. Dracula meets Wolfman: acceptance vs. partial belief. In *Induction, Acceptance, and Rational Belief*, ed. M. Swain, 157–85. Dordrecht: D. Reidel.

Joyce, J. 1998. A nonpragmatic vindication of probabilism. *Philosophy of Science* 65: 575–603.

Keiser, J. 2016. Bald-faced lies: how to make a move in a language game without making a move in a conversation. *Philosophical Studies* 173: 461–77.

Kelly, T. 2003. Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research* 66: 612–40.

Konek, J. and B. Levinstein. Forthcoming. The foundations of epistemic decision theory. *Mind*. doi: 10.1093/mind/fzw044.

Lackey, J. 2013. Lies and deception: an unhappy divorce. *Analysis* 73: 236–48.

Leland, P. 2015. Rational responsibility and the assertoric character of bald-faced lies. *Analysis* 75: 550–4.

Marsili, N. 2014. Lying as a scalar phenomenon: insincerity along the certainty-uncertainty continuum. In *Certainty-Uncertainty – and the Attitudinal Space in Between*, ed. S. Cantarini, W. Abraham and E. Leiss, 153–73. Philadelphia: John Benjamins.

Meibauer, J. 2005. Lying and falsely implicating. *Journal of Pragmatics* 37: 1373–99.

Michaelson, E. 2016. The lying test. *Mind and Language* 31: 470–99.

Ramsey, F.P. 1926/1990. Truth and probability. In *Philosophical Papers*, ed. D.H. Mellor, 52–109. Cambridge: Cambridge University Press.

Rees, C. 2014. Better lie! *Analysis* 74: 59–64.

Saul, J. 2012. *Lying, Misleading, and What is Said*. Oxford: Oxford University Press.

Schoenfield, M. 2015. Bridging rationality and accuracy. *Journal of Philosophy* 112: 633–57.

Schoenfield, M. Forthcoming. An accuracy based approach to higher order evidence. *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12329.

Sorensen, R. 2007. Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly* 88: 251–64.

Sorensen, R. 2010. Knowledge-lies. *Analysis* 70: 608–15.

Staffel, J. 2011. Reply to Roy Sorensen, 'Knowledge-lies'. *Analysis* 71: 300–2.

Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.

Stokke, A. 2013. Lying and asserting. *Journal of Philosophy* 110: 33–60.

Webber, J. 2013. Liar! *Analysis* 73: 651–9.