

Lying and Knowing*

Ben Holguín

August 2019

Abstract

This paper defends the simple view that in asserting that p , one lies iff one knows that p is false. Along the way it draws some morals about deception, knowledge, Gettier cases, belief, assertion, and the relationship between first- and higher-order norms.

1 Introduction

This paper argues for the following simple theory of lying:

KNOWING In asserting that p , S lies iff S knows that p is false.

KNOWING has a number of attractive features, many of which will emerge over the course of the discussion. For starters, it is a natural reflection of the standard dictionary definition of ‘lie’, which is ‘an intentionally false statement’.¹ For another, it provides a straightforward explanation of the core difference between lying and merely being wrong. When one is merely wrong, one doesn’t realize that what one says is false; when one lies, one does. This is presumably why when one is accused of lying, a plea of ignorance along the lines of ‘But I didn’t know what I was saying was false’ can be exculpating (if accepted).

Despite its *prima facie* plausibility, however, KNOWING has little support in the philosophical literature on lying. Aside from [Turri \(2016a\)](#) and [Benton \(2018b\)](#), I know of no proponents of the view. In fact, with the exception of those two I do not know of anyone who even *discusses* KNOWING, let alone endorses it.²

This paper argues that the relative obscurity of knowledge-theoretic accounts of lying is undeserved. Most of the paper plays defense: it explores general strategies the proponent of

*Forthcoming in *Synthese*. Thanks to Matthew Benton, Kyle Blumberg, Cian Dorr, Jane Friedman, Jeremy Goodman, Harvey Lederman, Andrew Lee, John Mullaney, Jake Nebel, Jim Pryor, Stephen Schiffer, Trevor Teitel, Peter Unger, and Jake Zuehl for helpful comments and discussion.

¹ See, e.g., <https://en.oxforddictionaries.com/definition/lie> as of mid 2019.

² Note, for instance, that KNOWING goes unmentioned in ([Mahon, 2016](#)), the *Stanford Encyclopedia of Philosophy*’s entry for ‘The Definition of Lying and Deception’, as well as in [Stokke’s \(2013b\)](#) survey article.

KNOWING can implement to defend her view against a variety of challenges. Some of these strategies have the form of a debunking explanation of recalcitrant intuitions, while others appeal to unorthodox but independently motivated stances one could take on the nature of knowing. But in general the idea is that the kinds of principles and examples that might lead one to reject KNOWING can in every case be accounted for on principled grounds. By way of making a positive case for the KNOWING, the paper argues that in addition to its *prima facie* plausibility, the view receives considerable support from issues concerning the norms of assertion. The idea, in short, is that if we think there are knowledge-theoretic norms of assertion, we should expect the concept of lying to be intimately connected with the concept of knowing.

2 The traditional view of lying

We'll start by comparing KNOWING with the traditional view on lying, whose proponents include Isenberg (1964); Chisholm & Feehan (1977); Kupfer (1982); Williams (2002); Frankfurt (2005); Mahon (2016). The traditional view takes the following two necessary conditions to be jointly sufficient for lying:

DECEPTION In asserting that *p*, *S* lies only if *S* intends to deceive her audience into believing that *p*.

BELIEF In asserting that *p*, *S* lies only if *S* believes that *p* is false.

That is to say: on the traditional view, lying is asserting what one believes to be false with the intention that others come to believe falsely as a result of it.

KNOWING departs from the traditional view in at least two respects. First, it does not entail DECEPTION. And second, due to the factivity of knowledge it (but not the traditional view) entails FALSITY:

FALSITY In asserting that *p*, *S* lies only if *p* is false.

With regards to BELIEF, KNOWING converges with the traditional view just to the extent we find correct the orthodox view that knowledge entails belief. And on that matter I am happy to assume it does.³

I will not defend KNOWING's departure from the traditional view on DECEPTION. There is already a sizable enough sub-literature on the case for and against DECEPTION—one that mostly sides against it—and for reasons of space it would be better not to recapitulate the dialectic here.⁴ Besides, DECEPTION could always be tacked onto KNOWING without

³ Though §6.2 will offer some tentative arguments for the view that knowing doesn't entail believing.

⁴ See, e.g., Carson (2006, 2010); Davis (1999); Fallis (2009, 2013, 2015); Sorensen (2007, 2010); Rutschmann & Wiegmann (2017) for arguments against DECEPTION, and see, e.g., Staffel (2011); Lackey (2013) for arguments in its defense.

compromising much of the view's theoretical interest. (In that case lying would be knowing that what you say is false with the intention of getting one's audience to believe it.)

The more interesting question (for our purposes) concerns FALSITY. Given that knowledge is factive, if lying requires knowing that what you assert is false, then lying requires asserting something false. Thus, if KNOWING is true, then contra the traditional view FALSITY has to be true too.

Like the issue of whether DECEPTION is true, there is a sizable sub-literature devoted to the issue of whether FALSITY is true. But unlike the view I take on DECEPTION, the view I take on FALSITY is not widely held. It has its proponents, but most in the literature seem to want to deny it.⁵

By and large the debate about FALSITY has been dominated by judgments about cases in which we see all the normal hallmarks of lying, except that the subject says something true. Many report the intuition that these are instances of lying. See for example the following case from Carson (2010, p. 16), which is supposed to be a paradigm example of the so-called *true lie*:

True lie I go fishing on a boat with a friend, John. He and I both catch a fish at the same time. Although we don't realize it, our lines are crossed. I have caught a very big fish and John has caught a little one, but we mistakenly believe that I caught the small fish and John caught the big one. We throw the two fish back into the water. I go home thinking that I caught a small fish. When I return, my father, an avid fisherman, asks me how I did. I say:

(1) I caught a very large fish and threw it back into the water.

thereby intending to deceive him about size of the fish that I caught.

Carson reports the intuition that (1) isn't a lie, but rather a case in which one tries to lie but fails. Jennifer Lackey (2013, fn. 9) reports the opposite: she takes **True lie** to "decisively refute" FALSITY. As far as I can tell, her intuitions are representative of most of those who work on lying.

Although I will not lean on intuitions about **True lie** much at all, I will report that I share Carson's judgments. I see it as no coincidence that the label 'true lie' sounds like an oxymoron.⁶ But rather than focus on judgments about cases like **True lie**—a practice that has not been particularly useful in settling things—I want to focus on judgments about certain kinds of sentences. In particular, I want to argue that if lying didn't entail falsity,

⁵ Carson (2006, 2010) is a vocal proponent of FALSITY. Turri & Turri (2015) provides empirical support in its favor, while Turri (2016a) and Benton (2018b) defend it indirectly by suggesting arguments in favor of KNOWING.

⁶ There is also some experimental evidence that these are the judgments of non-philosophers, at least once certain kinds of noise are controlled for: see (Turri & Turri, 2015). But see also (Wiegmann et al., 2016) for empirical evidence to the contrary.

then we'd expect there to be true readings of sentences like 'I know S spoke truly, but I don't know whether S lied', when in fact these sentences seem only to have false readings. A natural conclusion is that there is no such thing as a true lie.

Here are some examples of the sentences I have in mind:

- (2) ?? Although we know everything Bill Clinton is saying is true, what we don't know is whether he is lying.
- (3) ?? Even if what Lexie says is true, there is a good chance it is a lie.
- (4) ?? How many of the various true things Kyle told us were lies?

(2)–(4) all sound terrible. That fact is difficult to explain without recourse to *FALSITY*. Competent speakers are in general pretty good at imagining scenarios in which a sentence might be an appropriate thing to say; the fact that it is so difficult to do so for any of (2)–(4) suggests that the view that some lies express truths requires stronger arguments than those that are presently available.⁷ And against thinking that there is something peculiar about the choice of words, or that in general such sentences are hard to assess, consider analogous sentences but with 'mislead' or 'intend to mislead' in place of 'lie':

- (2*) Although we know everything Bill Clinton is saying is true, what we don't know is whether he is trying to be misleading.
- (3*) Even if what Lexie says is true, there is a good chance it is misleading.
- (4*) How many of the various true things Kyle told us were intended to be misleading?

Unlike (2)–(4), (2*)–(4*) all sound fine. This makes the explanatory challenge for the opponents of *FALSITY* presented by (2)–(4) even sharper.

Still: what of the intuitions of those like Lackey who take cases like **True lie** to show that one can assert something true and nonetheless lie? The hypothesis I like is that there is a tendency to slip between talk of lying and talk of trying to lie, and that this is what explains our judgments in cases of so-called true lies. Note that the phenomenon of conflating Φ ing with trying to Φ is not unique to lying: speeches like the following are in most contexts entirely appropriate:

⁷ In fact, I think similar considerations establish a *direct* argument in favor of *KNOWING*. It is highly unnatural to assert in one breath both that a person is lying and that they do not know whether what they say is true. Witness:

- (2_K) ? Although we know he is lying about whether he had an affair, what we don't know is whether he knows whether he had an affair.
- (3_K) ? Although for all Lexie knows everything she's saying is true; she is nonetheless lying through her teeth.
- (4_K) ? I would never lie *knowing* what I say is false; I only lie when there's a chance what I'm saying is true!

That these sentences seem bad is not at all easy to explain on the assumption that *KNOWING* is false.

- (5) Every morning I wake up and hit 1,000 serves. Only five were complete misses last time!
- (6) In a fit of passion Smith started shooting, but the gun was unloaded.

The lesson to take from examples is not that one can hit a serve while missing the ball, or that one can shoot a gun without firing any bullets; rather it's that for some verbs Φ , there are some contexts in which Φ is substitutable for 'try to Φ ' even when 'try to Φ ' fails to entail Φ . A natural hypothesis for the proponent of FALSITY to appeal to is that the same thing happens with 'lie'.⁸ And once one distinguishes carefully between speech acts that are lies and those that are merely attempted lies, the case in favor of the existence of true lies looks significantly less compelling.

I thus contend that the simplest account of the data is one that takes FALSITY to be a necessary condition on lying, and thus that the factivity of knowledge is a feature rather than a bug of KNOWING.

3 Lying, knowing, and believing

Supposing DECEPTION is false and FALSITY is true, the traditional view of lying is false. But that the traditional view is false is not much reason to think KNOWING is true. What it is reason to believe is that some instance of the following schema true (where Φ is a propositional attitude):

LYING SCHEMA In asserting that p, S lies iff: (i) S Φ s that p is false and (ii) p is false.

The remainder of the paper will take LYING SCHEMA for granted.⁹ The question is which instance of LYING SCHEMA we ought to accept. Here there are two candidates that suggest themselves: KNOWING, which is the view I ultimately defend, and BELIEVING, which is the traditional view minus DECEPTION and plus FALSITY:

⁸ This is not to give an explanation of *why* this phenomenon arises. But that task would take us beyond the scope of the present discussion. It suffices to know that, given our judgments about (2)–(4) *some* such debunking explanation of the intuitions in favor of the existence of true lies must be in order.

⁹ One objection to LYING SCHEMA whose discussion I will suppress concerns the status of the 'in asserting that p' bit of the schema. The objection arises from instances of metaphor, irony, and the like. To see why they might pose a problem, suppose for the moment that KNOWING is indeed the correct instance of LYING SCHEMA. On this view, in asserting that p one lies iff one knows that what one asserts is false. But of course when I assert something ironically or metaphorically I know (and *a fortiori* believe) that what I say is false. So aren't these cases counterexamples to LYING SCHEMA on just about any of its plausible implementations?

It is not entirely clear. Distinguish 'assert' as it is used in ordinary language from its technical uses in theorizing about speech acts, pragmatics, and epistemology. Supposing there is such a distinction—and I think that certainly ought to be the presumptive view—we can stipulate that the uses of 'assert' I am interested in are of the technical kind. On this understanding of 'assert', one can say 'p' without asserting that p. The natural move would then be to classify cases of irony and metaphor as cases in which one says 'p' without asserting that p. See (e.g.) Fallis (2009, 2010) and Stokke (2013a, 2018) for more discussion of these sorts of issues in the context of debates about lying.

KNOWING In asserting that p , S lies iff S knows that p is false.

BELIEVING In asserting that p , S lies iff: (i) S believes that p is false and (ii) p is false.

KNOWING eschews **LYING SCHEMA**'s factorization into two distinct conditions—the Φ condition and **FALSITY**—on account of the fact that the first condition entails the second when Φ is knowledge. Indeed, the fact that **KNOWING** reduces **LYING SCHEMA**'s two conditions to one is already some reason to prefer it to its competitors. If we have learned anything in the past half-century of epistemology, it's that philosophical analyses of epistemic concepts tend to be much more promising when stated in terms of knowledge than in terms of some combination of belief and truth. But high-altitude observations notwithstanding, the question of which of **KNOWING** or **BELIEVING** is the correct theory of lying should be settled by comparisons of explanatory adequacy.

The obvious place to start on that front is with an assessment of how the views compare in predicting our intuitive judgments about cases. And on that regard each of **KNOWING** and **BELIEVING** score equally well on the data considered so far—e.g., **True lie** and sentences like (2)-(4). Thus, different kinds of cases will be needed to gauge the relative (*prima facie*) empirical adequacy of the views.

The next two sections of the paper (§§4-5) will consider two empirical challenges against **KNOWING** and in favor of **BELIEVING**. The first challenge will focus on judgments about lying in Gettier cases, while the second will focus on judgments about lying in cases in which the speaker has strong probabilistic grounds for believing that what she asserts is false. I will argue that neither challenge is successful: in both cases the proponent of **KNOWING** has the resources to either deny or accommodate the putatively problematic judgments.

Before we get to those arguments, however, I want to make clear what their dialectical status is. We are assuming against the traditional view that **LYING SCHEMA** is true. What we are about to consider is whether there are empirical reasons to favor a belief-based implementation of the schema (**BELIEVING**) over a knowledge-based implementation (**KNOWING**). The aim of these next two sections is merely to show that certain considerations that might be taken to tell in favor of **BELIEVING** over **KNOWING** are in fact somewhat equivocal, and thus that broader theoretical considerations will be need to be called on to resolve the dispute. What might these broader considerations look like? Well, so far we've briefly mentioned the issue of elegance and simplicity. But it will not be until §6 that we see the main one: that we should expect there to be an intimate connection between lying and assertion, and that knowledge-based accounts of that connection are superior to belief-based ones. Until then the proponent of **KNOWING** will just be playing defense.

4 First challenge: Gettier lies

4.1 Gettier lies

The kinds of cases that make the sharpest trouble for KNOWING involve what I will call *Gettier lies*. These are cases in which a subject *S* asserts that *p* while being Gettiered on the proposition that *p* is false. That is to say: *S* has a justified true belief that $\neg p$ —the kind that in normal circumstances would lend itself to knowledge that $\neg p$ —and asserts that *p*. The issue is that the intuition that *S* lies is particularly strong in these cases, yet the view predicts otherwise. Here is an example:

Gettier lie Jones is driving through the country and absentmindedly forming beliefs of the form *there is a tree, there is a barn*, and so on. He decides to pull over to gaze upon a barn that is in front of him. While doing so, Jones receives a call from Smith, who asks him what he's doing. Jones ingeniously replies:

(7) I am not looking at a barn.

Jones speaks falsely: he is indeed looking at a barn. But Jones doesn't know that he speaks falsely: unbeknownst to him, he is in Fake Barn Country, a place littered with mules cleverly disguised as barns. He just happens to be looking at the one real barn in the area.

Does Jones lie? I imagine many readers will be of the mind that he does. After all, not only does Jones assert something false, he justifiably believes he asserts something false, and indeed he justifiably believes he *knows* he asserts something false. But supposing orthodoxy about knowledge in Gettier cases is correct, it follows from KNOWING that (7) can't be a lie. By contrast, BELIEVING counts (7) as a lie. So perhaps Gettier lies show that BELIEVING is the superior implementation of LYING SCHEMA.

The remainder of this section will argue against this conclusion: *pace* our intuitive judgments about cases like **Gettier lie**: we should either: (i) accept that the intuitions are mistaken and that Gettier lies aren't actually lies; or (ii) accept an interesting synthesis of KNOWING and BELIEVING: namely that in some (and perhaps all) contexts—but especially those in which issues concerning lies are at stake—knowledge just *is* having a (justified) true belief.¹⁰

4.2 First observation: true lies and Gettier lies

The case for the disjunction of (i) and (ii) draws on two observations.

¹⁰ Putting it this way is somewhat sloppy given the quantification over contexts. Really it should be something like: "In some contexts, 'knows' denotes the same relation as '(justifiably) believes truly'."

The first observation is that there is quite a lot in common between cases like **True lie** and **Gettier lie**. In fact, with the exception of the fact that one's assertion is true in the case of a "true lie" but false in the case of a Gettier lie, the speech acts are in the relevant *epistemic* respects identical. In **Gettier lie**, Jones (incorrectly) takes himself to know that he's looking at a barn, and tells Smith that he isn't. But the reason Jones is incorrect in taking himself to know is merely that he is Gettiered due to the presence of fake barns in his local environment. But suppose Jones happened to be in front of one of the many fake barns rather than the real one. Now when Jones asserts (7) ('I am not looking at a barn') he speaks truly: it is indeed the case that he is not looking at a barn. But on the epistemic side of things, there is no real difference. Jones just happens to be in one part of Fake Barn Country rather than another. From the inside, as it were, things are the same to Jones: in both cases he (incorrectly) takes himself to know that he is looking at a barn. It's just now he is incorrect because the belief is false rather than because the belief is Gettiered.

The fact that putative true lies and Gettier lies have this much in common is suggestive of a defensive strategy for the proponent of KNOWING. As we know from §2's discussion of FALSITY, there is an impulse to classify cases of true lies as lies. Indeed, many in the literature treat the impulse as a datum that a theory of lying ought to be designed to capture. But as we also know from §2, such an impulse is mistaken: if we know that someone's assertion is true, we know that assertion can't be a lie. Thus, since we know that we're already liable to make mistakes with 'lie' in cases like **True lie**, and since there are important underlying symmetries between the cases, it is reasonable to suspect that we might be doing the same thing in cases like **Gettier lie**.¹¹

Now, one might object to this line of argument on the following lines: though the relevant *epistemic* facts are identical between cases like **True lie** and **Gettier lie**, the *alethic* facts are not. So even if we sometimes confuse lies with merely attempted lies, we need independent argument to believe that it's the epistemic situation that matters when distinguishing between a lie and a mere attempt, rather than the alethic situation.

But this objection mistakes the dialectic. There is underdetermination of theory by the evidence. The evidence is that there is an impulse to call Gettier lies lies. The question is what our theory of lying should say about this impulse. One option is to treat it as a competent use of the concept, and thus to reject views like KNOWING. But another option is to look for a story that assimilates it with other cases in which the use of the concept is *not* competent. Supposing we have convinced ourselves that we are liable to mistake merely attempted lies for genuine lies (as in the case of true lies), then it is open to the proponent of KNOWING to claim Gettier lies are more of the same. Given the underlying epistemic symmetries between the cases, she has principled grounds on which to do so. Of

¹¹ Supposing we like the explanation of the mistake in the case of **True lie**, we could say the same of the mistake in the case of **Gettier lie**: that Jones' assertion of (7) is not a lie but a botched lie—i.e., a case of a subject trying but for purely external reasons failing to lie.

course, this does not mean we *ought* to follow the proponent of KNOWING here. That she has something to say about the problem doesn't mean she's right. But it does mean that a more sophisticated objection must be offered if one wants to convince oneself that she is wrong.

4.3 Second observation: Gettierized knowledge

The next argument against taking Gettier lies to refute KNOWING draws on observations familiar from the literature on contextualism about knowledge. It is that there are some contexts in which Gettierized true belief seems to suffice for knowledge: in particular, contexts in which the primary topic of inquiry is how the subject has acted or will act.¹²

Here is one example of the phenomenon from DeRose (2002, fn. 9) (citing lectures by Saul Kripke).¹³ When we're asking (as regards Watergate) 'What did Nixon know, and when did he know it?', we are unlikely to draw a distinction between the propositions Nixon knew and the propositions he merely truly believed. If Nixon was sure of something true, we will count him as having known it—regardless of whether he was Gettiered.

Here is another from Hawthorne (2000, p. 203), who puts the point clearly enough that it is best just to quote him:

A clock is stopped and reads "7 P.M." At 7 P.M., thinking that the clock is working, I set my watch by that clock. ...Much later that day (without having a chance to check my watch against other clocks in the interim), I look at my watch, see that it reads "10 P.M." and think to myself "It's 10 P.M.; The party begins at 10:30, so it's time to leave for the party." Someone asks you whether I know that the party begins in half an hour. You are fully apprised of the Gettierized etiology of the relevant temporal beliefs. What do you say? Once again, nearly everyone has a strong pull in this context to say "Yes" and thus to concede that in this case we have Gettierized belief that counts as knowledge.

Call the examples from DeRose and Hawthorne instances of *Gettierized knowledge*. From the perspective of a theory of knowledge, what are we to make of the fact that the ordinary usage of 'know' seems to admit these readings? I see three options: (i) take Gettierized knowledge to show that knowing *just is* believing truly; (ii) take our judgments about (so-called) Gettierized knowledge to be in error, and look for some non-semantic story to account for the data; or (iii) take Gettierized knowledge to show that 'know' is context-sensitive, and that in at least some contexts it denotes the same relation as the one denoted by 'believe truly'.

¹² See, e.g. (DeRose, 2002, fn. 9) for discussion of some of the cases that motivate this conclusion (including the one from Hawthorne). See also (e.g.) (Turri, 2016b, 2017).

¹³ The Kripke lectures DeRose cites eventually appeared in print as (Kripke, 2011), though unfortunately the print version omits the relevant example.

Although the question of which of (i)-(iii) to prefer is interesting and important, it is also irrelevant to the question of the nature of lying. For as I will now argue, regardless of whether we prefer (i), (ii), or (iii), the fact that the *disjunction* of (i)-(iii) is true puts KNOWING in no worse a position than BELIEVING with regards to cases like **Gettier lie**.

The argument is by cases. Start with (i), the view that the lesson to take from cases of Gettierized knowledge is that knowledge just is believing truly. This case is the easiest, as it involves the collapse of KNOWING and BELIEVING. If knowing *just is* believing truly, then the view that says that one lies in asserting *p* iff one knows that *p* is false *just is* the view that one lies in asserting *p* iff *p* is false and one believes that *p* is false. On this way of thinking about things, **Gettier lie** is no counterexample to KNOWING for the simple reason that KNOWING (like BELIEVING) predicts that it is a case of lying.

Now consider (ii), the view that the conventional wisdom about knowledge is correct: in no context are ‘know’ and ‘believes truly’ coextensive; rather, what explains the judgments to the contrary is an error theory about our inclination to conflate the two. What might such an error theory look like? A full discussion of that question is a topic for another paper, but for present purposes it is enough to draw attention to a feature in common to just about every putative instance of Gettierized knowledge. Put roughly, it is that each of these cases is associated with a context in which the topic of discussion concerns how an agent has or will *act*. In the Nixon case, our concern is with assessing what may have motivated his decisions surrounding Watergate, while in the clock case, our concern is assessing the subject’s capacity to make a party on time. Given the existence of a unifying feature, one would naturally expect it to make an appearance in the error theory. A toy version might thus look something like: when our interest is in assessing what motivates *S* to act, we sloppily fail to distinguish *S*’s truly believing that *p* and *S*’s knowing that *p*.

Whatever its exact details, it is plausible that an error theory along these lines may be used by the proponent of KNOWING to explain away the judgment that Gettier lies are lies. Take **Gettier lie**. Jones believes that he is looking at a barn, and indeed he *is* looking at a barn. But unbeknownst to Jones, it’s the only real barn in the area; the rest are mules cleverly disguised as barns. So (according to the view in question) Jones can’t know that he is looking at a barn. It thus follows from KNOWING that Jones’s assertion of (7) (‘I am not looking at a barn’) cannot be a lie. However, as we know from our error theory, there is sometimes a temptation to conflate knowing and believing truly. In particular, we tend to conflate the two when the topic of discussion concerns something like the subject’s motivation for her actions. This is exactly the sort of context we should expect to often be associated with assessments of lying. Indeed, given the strength of the intuitions in the Nixon example—a paradigm (putative) example of Gettierized knowledge—it is plausible that the inclination to conflate true belief with knowledge is more pressing when the *moral* qualities of *S*’s actions are salient. This adds to the case that whatever is driving our judg-

ments about **Gettier lie** might not be semantically respectable. Thus, to the extent that there is a more general phenomenon whereby we find ourselves confused about whether a subject knows versus happens to believe truly, we have reason to think that our judgments about cases like **Gettier lie** are just more of the same. Jones doesn't know and thus can't be lying, but because he believes truly we are mistakenly inclined to treat him as one who knows, and thus as one who lies.

Finally consider (iii), the view that cases of Gettier knowledge show that the proper semantics for 'know' is contextualist, and that at least in some contexts 'know' and 'believes truly' are coextensive. Supposing this view is correct, 'know' denotes a multitude of relations. So if lying is to be understood in terms of knowing, 'lie' must too. This raises the question of how to properly formulate a principle like KNOWING. The straightforward option is to go meta-linguistic:

META-LINGUISTIC KNOWING In every context *c*: If *S* asserts that *p*, then: 'S lies in asserting *p*' is true in *c* iff 'S knows that *p* is false' is true in *c*.

The natural hypothesis for proponents of META-LINGUISTIC KNOWING is that our anti-KNOWING intuitions are due to equivocation: in the case of **Gettier lie**, when assessing 'Jones lies', the contexts naturally invoked are those in which 'know' is coextensive with 'believes truly'; but when assessing 'Jones knows', the contexts naturally invoked are those in which 'know' is *not* coextensive with 'believes truly'. That is to say, when one thinks about lying, one will be more inclined to treat on a par knowing and believing truly; but when one thinks about knowledge, one will be more inclined to treat them differently. There is thus a true reading of 'Jones lies' and a true reading of 'Jones doesn't know that what he asserts is false', but no *uniform* true reading of 'Jones lies even though he doesn't know that what he asserts is false'. Since META-LINGUISTIC KNOWING requires this sort of uniformity, **Gettier lie** is no counterexample to KNOWING on its intended interpretation.

The lesson, then, is that no matter which of (i), (ii), or (iii) is true, the fact that one of them has to be is enough to establish that BELIEVING has no clear empirical advantage over KNOWING on the topic of Gettier lies.

We thus have two independent bits of data that suggest that *if* there is reason to prefer BELIEVING to KNOWING, it is not because of our intuitions about Gettier lies: first, the unity of Gettier lies and so-called true lies; and second, the unity of the contexts in which we are inclined to call Gettier lies lies and the contexts in which we are inclined to count Gettierized true belief as knowledge.

5 Second challenge: lottery lies

5.1 Lottery lies

That covers the first challenge to KNOWING in favor of BELIEVING. Before presenting the second challenge, it is worth making an observation about the kind of belief involved in Gettier cases. When one is Gettiered with regards to one's belief that $\neg p$, one (falsely) takes oneself to know that $\neg p$. As such, the kind of belief involved in Gettier cases is a paradigmatic instance of *full* or *outright* belief: a kind of doxastic attitude that goes beyond merely finding the proposition in question probable.

But of course it is possible to falsely assert that p even when one's attitude toward $\neg p$ is weaker than full belief. This will happen whenever (i) p is false and (ii) one's evidence for $\neg p$ is merely statistical. For example:

Lottery Jones knows that yesterday a fair 100-sided die was rolled, but he has no idea how it landed. Smith wrongly thinks that Jones knows how it landed, so Smith calls Jones up and asks him about it. Knowing that Smith wrongly thinks Jones knows, Jones replies:

(8) The die landed 100.

Jones speaks falsely: the die in fact landed 13.

Ought we think that assertions like (8) are lies too?

Although I expect opinions to be less clear here than they are in the case of **Gettier lie**, I myself don't find it very difficult to get into a frame of mind in which it would be fair to say that Jones lies. Again, he asserts something false, and he is well aware that there is a 99% chance that what he asserts is false. So let us suppose for now that assertions like (8) (henceforth *lottery lies*) are in fact lies. What does this tell us about KNOWING and BELIEVING?

Well, assuming (as is standard) that one cannot know on purely statistical grounds that a fair 100-sided die didn't land 100, (8) is a counterexample to KNOWING. Whether it is a counterexample to BELIEVING depends on whether it is possible to believe that p without fully believing it. If it *isn't* possible to believe that p without fully believing that p , then (8) is a counterexample to BELIEVING too, for Jones would not believe that what he asserts is false (rather, he'd merely believe it 99% likely). But if it is possible to believe without believing fully—say because believing is just a matter of believing sufficiently likely—then BELIEVING is left unscathed.

There is an enormous literature on the question of whether the attitudes of belief and full belief are identical, and for reasons of space we do better leaving it alone.¹⁴ So let

¹⁴ For examples of views that defend the view that believing is fully believing (rather than just believing sufficiently

us bypass the issue by temporarily suspending talk of believing, and instead speak directly about believing fully and believing likely. We will leave open the exact threshold needed for a proposition to be believed likely (as well as, e.g., the question of whether the threshold might vary across contexts). We will only assume (i) that it is possible to believe likely on purely statistical grounds and (ii) that believing *p* likely need not entail believing *p* 100% likely. This gives us two implementations of the LYING SCHEMA:

BELIEVING FULLY In asserting that *p*, *S* lies iff: (i) *S* believes fully that *p* is false and (ii) *p* is false.

BELIEVING LIKELY In asserting that *p*, *S* lies iff: (i) *S* believes likely that *p* is false and (ii) *p* is false.¹⁵

So: if lottery lies like (8) really are lies, then **Lottery** tells against both **KNOWING** and **BELIEVING FULLY** and in favor of **BELIEVING LIKELY**. It is thus incumbent on both the proponents of **KNOWING** and **BELIEVING FULLY** to provide an argument against treating lottery lies as the genuine article. We now turn to offering one.

5.2 Bullshit

It is uncontroversial that there is something wrong with asserting something false. But it is also uncontroversial that what's wrong with it needn't always be that it is dishonest. If one takes oneself to know that *p* and asserts *p* sincerely, then one's assertion that *p* is not dishonest.¹⁶ But if one asserts that *p* knowing full well that one doesn't know that *p*, then one's assertion is dishonest. And that's because an assertion that *p* typically implies that one knows that *p*. How does this implication arise? Well, we'll have more to say about that in §6, but for now it suffices to a platitude about the practice of assertion: in ordinary contexts, when a speaker asserts that *p*, she can expect her audience to infer that she knows that *p*.¹⁷ Thus, if I assert that *p* I can expect my audience to infer that I know that *p*. It follows that if I know that I don't know that *p*, then I know that by asserting *p* I will cause my audience to believe something false—namely, that I know that *p*. That is dishonest.

There is thus an important difference between assertions *knowingly* made in ignorance from those *unknowingly* made in ignorance. We may capture that difference with the following (somewhat) stipulative definition of *bullshit*:

likely), see, e.g., Hintikka (1962); Stalnaker (1984); Williamson (2000); Buchak (2014); Ross & Schroeder (2014); Staffel (2016). For examples of views that entail that believing is just a matter of believing sufficiently likely, see, e.g., Sturgeon (2008); Foley (2009); Leitgeb (2013); Hawthorne *et al.* (2016); Dorst (2019).

¹⁵ Views closely resembling **BELIEVING LIKELY** have been defended in the literature by Marsili (2014) and Krauss (2017). See (Benton, 2018a) for criticisms of Krauss' view.

¹⁶ Well, perhaps it can be dishonest by being pragmatically misleading. We can suppress such complications here.

¹⁷ Note: this is *not* the claim that there are constitutive connections between knowledge and assertion. It's just the claim that people ordinarily make this inference and know that others ordinarily make it too.

BULLSHIT In asserting that p , S bullshits iff S knows that S does not know whether p .¹⁸

What's the relevance of bullshit to so-called lottery lies? Well, assertions like (8) are definitely bullshit: Jones knows that he doesn't know whether the die landed 100, yet he asserts that it did. And we know that bullshit is dishonest. Thus, a hypothesis presents itself: the temptation to call an assertion like (8) a lie is due to a tendency to mistake bullshit for lying.

The bullshit hypothesis gives the proponent of KNOWING (as well as BELIEVING FULLY) a principled response to the threat from lottery lies. Ought it be the default hypothesis? Maybe not. But given that we know that intuitions about lying are messy and perhaps even sometimes systematically mistaken, and given that it is not even clear how sharp the intuition that (8) is a lie is, it is certainly a reasonable one to appeal to. Moreover, if we have independent reasons to doubt the probative force of intuitions that compel us to classify lottery lies as genuine lies, then we should be more inclined to take the bullshit hypothesis seriously. I will now give a brief argument in favor of this line of thought.

5.3 Casting doubt on lottery lies

By my lights, the central reason to doubt the probative force of our intuitions about lottery lies is fairly unsophisticated: it is that the intuitions are just too messy to be given much weight.

To help bring out the messiness, consider a case like **Lottery**, but in which a speaker takes the odds that what they assert is false to be approximately even:

Fair coin Jones knows that yesterday a fair coin was flipped, but he has no idea how it landed. Smith wrongly thinks that Jones knows how it landed, so Smith calls Jones up and asks him about it. Knowing that Smith wrongly thinks Jones knows, Jones replies:

(9) The coin landed heads.

Jones speaks falsely: the coin in fact landed tails.

Does Jones lie? Well, on the one hand it does not seem particularly difficult to get into a frame of mind in which, upon discovering that the coin landed tails and that Jones knew he had no idea how it landed, one acts appropriately in accusing Jones as having lied. But on the other hand it also does not seem particularly difficult to get into a frame of mind

¹⁸ See Frankfurt (2005) for a classic discussion on bullshit. See also Fallis & Stokke (2017) for a discussion of the relationship between lying and bullshit. As I am using 'bullshit', much of what I say here may be in tension with what Fallis and Stokke want to say about it. But since I am using the term in something of a stipulative sense, we needn't worry about comparing the views. For those who have strong views about how the word 'bullshit' ought to be used, I'll note that the same points could be made with 'making things up' in its place.

in which it would be better to accuse Jones merely of bullshitting, having made things up, etc. So though I lean strongly toward a negative answer, without recourse to theoretical considerations it is not entirely obvious what to say.

Regardless of the initial impressions, however, I think we should be confident that no good theory of lying will classify assertions such as (9) as lies. By extension, we should be suspicious of whatever mechanism is driving us to classify them that way. Jones doesn't even believe it *likely* that what he says is false. And if believing one's assertion likely to be false isn't a necessary condition on lying, then what is? Merely having *some* confidence that what one says is false? Absent further constraints, such a view would have it that if, on purely statistical grounds, I falsely assert of a 100-sided die that it didn't land 100—a proposition which I have very good reason to suspect true—I happen to be lying. That is not a promising result. And though it may be possible to strain hard enough to get 'lie' to seem like an appropriate word to describe such an assertion, following such impulses wherever they lead seems to me bad methodology. The underlying intuitions are too messy and the theories to which they lead are too *prima facie* implausible to reactively adhere to them. This is especially true once we notice the availability of alternative diagnoses of the felt impropriety of Jones' speech act—in particular, a diagnosis in terms of the fact that Jones is asserting something he knows he doesn't know to be true.¹⁹ As such, the lesson to take from examples like **Coin flip** is that we should be generally suspicious of our judgments about lying in cases of assertions made on purely statistical grounds. This is not to say we should discount them entirely. But we should recognize that they are too equivocal to be treated as an independent test on the empirical adequacy of a theory of lying.²⁰

¹⁹ One might complain that this argument is soritical, as I'm inferring from the fact that **Fair coin** isn't a case of lying that a case that's sufficiently like it (**Lottery**) can't be either. Perhaps what is "really" revealed by our pattern of judgments across **Lottery** and **Fair coin** is that part of the vagueness of 'lie' is in terms of the degree of confidence one must have in the falsity of what one asserts to count as lying. On this view **Lottery** is a case of lying and **Fair coin** is not; the border between a lie and mere dishonesty falls somewhere between the two cases; and either way both (and all the possible cases between) involve bullshit.

I have no direct objections against the view that 'lie' is vague in this way. Like anything, it is an option to be explored. But the charge of soriticality is mistaken. All I have claimed is that there is a certain status in common to the assertions in both **Fair coin** and **Lottery**—one we know deserves the dis-honoric 'dishonest'—and that the residual differences in the status of those assertions have the feel of being degree-like rather than kind-like. Nothing about that line of argument is soritical. Whether or not it is correct, the hypothesis that these cases are examples of bullshit rather than lies has a good amount of explanatory power. And given the unclarity in the judgments about these thought experiments, the best bet towards a theory of lying will require lots of abductive inference of this sort.

²⁰ Though see, e.g., (Marsili, 2014) and (Krauss, 2017) for more optimistic takes on this methodological stance. I imagine such authors would claim a theoretically interesting difference in the strength of their intuitions about **Lottery** and **Fair coin**. I myself do not detect one.

6 Lying and the norms of assertion

We've now considered and rejected two arguments in favor of BELIEVING (or at least BELIEVING LIKELY) over KNOWING: first (§4), an argument from so-called Gettier lies; and second (§5), an argument from so-called lottery lies. On all the other conceptual and empirical considerations—the distinction between lying and merely being wrong, so-called “true lies”, etc.—the views seem to score about as well as each other. What, then, is supposed to make the case for KNOWING over BELIEVING?

Aside from considerations of simplicity and elegance, I think the best case for KNOWING over BELIEVING comes out of considerations of the relationship between the norms of assertion and lying. Here is the idea in broad outline.²¹ Supposing there is a norm that says that one may assert that p only if one knows that p , we should expect there to be something particularly bad about what I will call *anti-assertion*: the act of asserting the opposite of what you know to be true. Moreover, given the special standing of this category of assertion, it would be entirely unsurprising if natural language offered a device for talking about it. An attractive hypothesis is thus that lies are anti-assertions. The badness of other violations of the norm of assertion can then be explained derivatively, and in a way that manages to account for some of the impulse to categorize merely intended lying and bullshit as lying proper. Thus, to the extent that we prefer knowledge-theoretic accounts of the norms of assertion to belief-theoretic accounts, we should think KNOWING is the correct theory of lying.

6.1 A taxonomy of assertions

Let us suppose as is widely accepted that *some* instance of the following schema is true (where Φ denotes an attitude):²²

ASSERTION SCHEMA S may assert that p only if $S \Phi$ s that p .²³

Given ASSERTION SCHEMA, the following (non-exhaustive) taxonomy of assertion suggests itself:

- First-order categories for S 's assertion that p :

$S \Phi$ s THAT P : Proper assertion.

$S \Phi$ s THAT $\neg P$: Anti-assertion.

²¹ See Benton (2018b) for a similar exploration of the relationship between lying and the norms of assertion.

²² Weiner (2005) is a notable exception: he defends the norm that S may: assert that p iff p . See (Weiner, 2007) and (Pagin, 2016) for overviews on the question of LYING SCHEMA. See also (Cappelen, 2011) for skepticism about the general project of theorizing about the rules of assertion.

²³ The force of the ‘may’ should be understood purely epistemically. Assertions that are improper because they are rude, pragmatically misleading, etc., are to be ruled out on other grounds.

S DOESN'T Φ WHETHER P: Sloppy assertion.

- Second-order categories for S's assertion that p:

S FALSELY Ψ S THAT S Φ S THAT P: Botched proper assertion.

S FALSELY Ψ S THAT S Φ S THAT \neg P: Botched anti-assertion.

S Ψ S THAT: S NEITHER Φ S THAT P NOR Φ S THAT \neg P: Reckless assertion.

For present purposes I will stay neutral on whether and when Φ and Ψ denote different attitudes. But if one is attracted to an implementation of ASSERTION SCHEMA in terms of knowledge (which I will defend in a moment), then the attitudes must be distinct for botched proper assertion and botched anti-assertion. (You can't falsely know something.) Φ would be knowing, and Ψ would presumably be something like believing.

With this taxonomy in hand, we can begin to tell a natural story of what grounds the felt impropriety of the different ways of failing to assert properly:²⁴

- Anti-assertion: One anti-asserts when one is in a position to assert in accordance with the norm on assertion, and instead one asserts the negation of what one ought to have asserted.²⁵ In anti-asserting one thus misleads doubly: not only does one assert what one is prohibited from asserting, one *could have* been cooperative and asserted something informative.
- Sloppy assertion: Sloppy assertion is another way to fail to live up to the norm of assertion, but is bad *per se* only to the extent that failure to live up to the norm of assertion is bad *per se*. As we will see shortly, the extent of that impropriety can vary quite strongly in accordance with one's second-order attitudes toward the proposition one sloppily asserts.
- Botched proper assertion: This is defective assertion with an excuse (and perhaps a justification). By the standards of the norm of assertion (i.e. Φ), you took yourself to be living up to the norm of assertion. Plausibly, then, botched proper assertion is internally indistinguishable from proper assertion, and thus carries with it all the exculpating benefits of internal indistinguishability.
- Botched anti-assertion: Just as botched proper assertion is internally indistinguishable from proper assertion, botched anti-assertion is internally indistinguishable from anti-assertion. Roughly speaking then, when one botches an anti-assertion, it is only

²⁴ (Benton, 2018b, §9.4) features a similar discussion of first- and second-order norms of assertion. He offers the categories of "reasonable assertion" (believing that you know that p), "negligent assertion" (not considering whether you know that p) and "vicious assertion" (believing that you know that \neg p). I leave the questions of the relationship between these categories and my own (none is a perfect match)—as well as their usefulness in accounting for our judgments about lies, bullshit, and other speech acts—for future work.

²⁵ Or something that entails its negation.

because one's environment is uncooperative that one doesn't manage to produce an anti-assertion. Thus, given a weak normative anti-luck principle like:

NORMATIVE ANTI-LUCK If as a result of Xing S_B is (or would be) subject to normative scrutiny, then for any other person S_C that is in the relevant respects an internal duplicate of S_B : S_C is subject to some amount of normative scrutiny too.

—we have a straightforward explanation of why botched anti-assertions seem wrong in basically the same way actual anti-assertions seem wrong.

- Reckless assertion: Reckless assertion is bad for whatever reasons deliberately violating the norm on assertion is bad. Here's a stab at it: given that speakers normally assume that other speakers are abiding by the norm of assertion, in recklessly asserting that p one Ψ ingly misleads one's audience into thinking that one Φ s that p . This is dishonest. What the relation is between this kind of dishonesty and the dishonesty of anti-assertion is a question for other work.²⁶ Still, even without a sharp sense of what the difference amounts to, it is plausible that the categories of anti-assertion and reckless assertion should be felt to be of different statuses.

Now for the official view: Anti-assertion—and only anti-assertion—is lying. Reckless assertion is bullshitting. Botched anti-assertion—of which both **True lie** and **Gettier lie** are instances—is merely attempted lying. All involve dishonesty, but not in the same way. When one lies, one asserts the opposite of what one would assert if one were abiding by the norm of assertion. Not only does one lead one's audience away from the truth, but one could have led them to it had one wanted to. We have reason to expect special condemnation for this kind of speech act (at least in normal cases), so it is no surprise that we would have a concept that tracks it. When one bullshits, by contrast, one doesn't Ψ ingly lead one's audience away from the truth about what one asserts. Still, one Ψ ingly misleads one's audience into thinking that one Φ s the content of one asserts, thereby encouraging one's audience to accept the content of what one asserts without the appropriate kind of evidence for it. But it's not like one could have abided by the norm of assertion and asserted the truth of the matter had wanted to: if one Ψ s that one doesn't Φ *whether* p , then one Ψ s that one doesn't Φ *that* p . Finally, when one botches a lie in asserting that p , one does something very much like what one does when successfully lies in asserting that p . It's just that—unlike in the case of actual lying—when one botches it it's because one wasn't in a position to assert that p after all. Sometimes one will botch it by asserting the truth (as in **True lie**), other times one will botch it by asserting a falsehood that one doesn't Φ (as in **Gettier lie**). But in neither case does one lie, as again one couldn't have properly asserted

²⁶ On the assumption that anti-assertion is lying and reckless assertion is bullshit, (Fallis & Stokke, 2017) is a good place to start.

the truth had one wanted to. But given *NORMATIVE ANTI-LUCK*, one does something that in many of the normative respects is as bad as lying.

The view that lying is anti-assertion has an attractive amount of explanatory power. When combined with some observations about the various ways in which one can fail to abide by the primary norm of assertion, it works to account for a whole host of messy data—including, most importantly, our impulse to treat similarly lies, bullshit, and merely attempted lies.

6.2 The knowledge norm on assertion

We now have the resources to offer a concrete argument in favor of *KNOWING* over *BELIEVING*. It is that we should prefer *KNOWING* to *BELIEVING* because we should prefer *KNA* to *BNA*:

KNA S may assert that p only if S knows that p.

BNA S may assert that p only if: (i) S believes that p and (ii) p is true.

On the issue of why *KNA* is a better norm than *BNA*, I do not intend to add much to the already quite sizable literature in defense of *KNA* over other rival norms.²⁷ But here are two considerations that seem to me among some of the most decisive reasons to prefer *KNA* to *BNA*.

First, certain Gettier-like cases present evidence that one's truly believing that p is insufficient grounds for the corresponding assertion. Consider, e.g.:

Crystal ball Jones has been told by highly reliable sources that his crystal ball is a highly reliable source of information about the local weather. He checks it for tomorrow's forecast and it tells him that it is going to rain. Taking himself to know that it will rain tomorrow, Jones tells Smith:

(10) It will rain tomorrow.

In fact Jones speaks truly: it will rain tomorrow. But the crystal ball is broken, and in fact always says it will rain.

Suppose Smith knows the etiology of Jones's belief, as well as the fact that it will indeed rain tomorrow. Is Smith in a position to complain about Jones's assertion? It is intuitive to think that he is. He may rightly complain that although what Jones says is true, and that although Jones believes that what he says is true, since Jones doesn't *know* that what he is saying is true he shouldn't assert it. That is to say, Smith may rightfully respond along the

²⁷ See, e.g., Unger (1975); Williamson (2000); DeRose (2002); Hawthorne (2004); Fantl & McGrath (2009); Benton (2011); Turri (2011); Worsnip (2017).

lines of ‘Actually, you don’t know that: contrary to what you’ve been told, your crystal ball is not a useful guide to the local weather’.

Second, there is reason to believe that knowing need not entail believing. The kinds of cases that come to mind here are those first brought to the scene by [Radford \(1966\)](#):²⁸

Unconfident examinee Jane feels quite certain that she does not know any English history. But Jane has studied quite a lot, and so whenever she is asked to provide dates for certain events in English history, she does so correctly. She just happens to feel like she is guessing. Smith asks Jane the year of Queen Elizabeth’s death. Jane unconfidently replies:

(11) Queen Elizabeth died in 1603.

Jane is correct: Queen Elizabeth died in 1603. Moreover, the causal etiology of Jane’s is the sort that normally produces knowledge.

Intuitively, Jane knows that Queen Elizabeth died in 1603. Less but (perhaps) still reasonably intuitively, Jane does not believe that Queen Elizabeth died in 1603. **Unconfident examinee** is thus a putative example of knowledge without belief. Supposing the example is genuine, what are we to make of Jane’s assertion of (11)?

My intuition is that the assertion, though not fully ideal (for reasons to be explained in a moment), is nonetheless appropriate in the ways that matter. To the extent we are in the frame of mind in which Jane really does know what she asserts, any *direct* challenge should seem odd.

Still, Jane’s speech act is not unimpeachable. And that’s because in virtue of failing to believe that Queen Elizabeth died in 1603, Jane will (if rational) fail to believe that she knows it too. Jane’s assertion of (11) is thus a case in which an agent asserts in accordance with the norm of assertion while thinking she fails to assert in accordance with the norm of assertion. So long as we think it matters that one not violate the norm of assertion—and again we needn’t have a settled view on what makes such a thing matter—we should expect it to matter derivatively that one not do something which one *thinks* will violate the norm of assertion.

So supposing KNA is correct, we predict that Jane does do something wrong: she asserts (11) when she thinks doing so will bring about a violation of the norm of assertion. That prediction seems to me a good one: it is in general reasonable to complain when one asserts what one isn’t sure of—see, e.g., the practice of asking ‘Are you sure?’ in response to assertions of the relevant sort. But the proponent of KNA has the resources to explain

²⁸ My presentation of the case is borrowed from [Myers-Schulz & Schwitzgebel \(2013, p. 372\)](#), who discuss empirical findings related to cases like Radford’s. See also [Rose & Schaffer \(2013\)](#) for a critical response.

all on her own terms. And given our intuitions about assertions in cases like **Unconfident examinee**, the story seems to strike the right balance.

By contrast, the proponent of BNA must come down more strongly on Jane's assertion of (11). On that view, Jane's assertion is straightforwardly defective. To account for its felt *appropriateness*, then, the proponent of BNA must either derive or posit some norm that explains why we tend to let people off the hook when they assert what they merely know. But if we take the judgments about **Unconfident examinee** at face value, then knowing entails neither believing nor believing that one knows. So it is unclear how such a norm is supposed to be derived. And positing it by fiat seems uncomfortably ad hoc, at least in comparison to the more unified story the proponent of KNA has to offer. Cases like **Unconfident examinee** thus provide compelling evidence in favor of knowledge-theoretic accounts of the norm of assertion over belief-theoretic accounts.

This gives us the argument: since KNA is the true instance of the ASSERTION SCHEMA, and since lying is anti-assertion, KNOWING must be the true instance of the LYING SCHEMA. To lie is to know that what you say is not so. Everything else is botched lying or bullshit.

References

- Benton, Matthew A. 2011. Two More for the Knowledge Account of Assertion. *Analysis*, **71**(4), 684–687.
- Benton, Matthew A. 2018a. Lying, Accuracy and Credence. *Analysis*, **78**(2), 195–198.
- Benton, Matthew A. 2018b. Lying, Belief, and Knowledge. In: Meibauer, Jörg (ed), *The Oxford Handbook of Lying*. Oxford University Press.
- Buchak, Lara. 2014. Belief, Credence, and Norms. *Philosophical Studies*, **169**(2), 1–27.
- Cappelen, Herman. 2011. Against Assertion. In: Brown, Jessica, & Cappelen, Herman (eds), *Assertion: New Philosophical Essays*. Oxford University Press.
- Carson, Thomas L. 2006. The Definition of Lying. *Noûs*, **40**(2), 284–306.
- Carson, Thomas L. 2010. *Lying and Deception: Theory and Practice*. Oxford University Press.
- Chisholm, Roderick M., & Feehan, Thomas D. 1977. The Intent to Deceive. *Journal of Philosophy*, **74**(3), 143–159.
- Davis, Wayne A. 1999. Communicating, Telling, and Informing. *Philosophical Inquiry*, **21**(1), 21–43.
- DeRose, Keith. 2002. Assertion, Knowledge, and Context. *Philosophical Review*, **111**(2), 167–203.
- Dorst, Kevin. 2019. Lockeans Maximize Expected Accuracy. *Mind*, 1–37.
- Fallis, Don. 2009. What Is Lying? *Journal of Philosophy*, **106**(1), 29–56.
- Fallis, Don. 2010. Lying and Deception. *Philosophers' Imprint*, **10**.
- Fallis, Don. 2013. Davidson Was Almost Right About Lying. *Australasian Journal of Philosophy*, **91**(2), 337–353.
- Fallis, Don. 2015. Are BaldFaced Lies Deceptive After All? *Ratio*, **28**(1), 81–96.
- Fallis, Don, & Stokke, Andreas. 2017. Bullshitting, Lying, and Indifference Toward Truth. *Ergo: An Open Access Journal of Philosophy*, **4**.
- Fantl, Jeremy, & McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- Foley, Richard. 2009. Beliefs, Degrees of Belief, and the Lockean Thesis. Pages 37–47 of: Huber, Franz, & Schmidt-Petri, Christoph (eds), *Degrees of Belief*. Springer.
- Frankfurt, Harry G. 2005. *On Bullshit*. Princeton University Press.
- Hawthorne, John. 2000. Implicit Belief and A Priori Knowledge. *Southern Journal of Philosophy*, **38**(S1), 191–210.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, John, Rothschild, Daniel, & Spectre, Levi. 2016. Belief is Weak. *Philosophical Studies*, **173**(5), 1393–1404.
- Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- Isenberg, Arnold. 1964. Deontology and the Ethics of Lying. *Philosophy and Phenomenological Research*, **24**(4), 463–480.

- Krauss, Sam Fox. 2017. Lying, Risk and Accuracy. *Analysis*, 77(4), 726–734.
- Kripke, Saul A. 2011. Nozick on Knowledge. In: Kripke, Saul A. (ed), *Philosophical Troubles. Collected Papers Vol I*. Oxford University Press.
- Kupfer, Joseph. 1982. The Moral Presumption Against Lying. *Review of Metaphysics*, 36(1), 103–126.
- Lackey, Jennifer. 2013. Lies and Deception: An Unhappy Divorce. *Analysis*, 73(2), 236–248.
- Leitgeb, Hannes. 2013. Reducing Belief Simpliciter to Degrees of Belief. *Annals of Pure and Applied Logic*, 164(12), 1338–1389.
- Mahon, James Edwin. 2016. The Definition of Lying and Deception. In: Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, winter 2016 edn. Metaphysics Research Lab, Stanford University.
- Marsili, Neri. 2014. Lying as a Scalar Phenomenon. In: Cantarini, Sibilla, Abraham, Werner, & Leiss, Elizabeth (eds), "*Certainty-uncertainty – and the attitudinal space in between*",. John Benjamins Publishing.
- Myers-Schulz, Blake, & Schwitzgebel, Eric. 2013. Knowing That P Without Believing That P. *Noûs*, 47(2), 371–384.
- Pagin, Peter. 2016. Assertion. In: Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, winter 2016 edn. Metaphysics Research Lab, Stanford University.
- Radford, Colin. 1966. Knowledge—By Examples. *Analysis*, 27(1), 1–11.
- Rose, David, & Schaffer, Jonathan. 2013. Knowledge Entails Dispositional Belief. *Philosophical Studies*, 166(1), 19–50.
- Ross, Jacob, & Schroeder, Mark. 2014. Belief, Credence, and Pragmatic Encroachment¹. *Philosophy and Phenomenological Research*, 88(2), 259–288.
- Rutschmann, Ronja, & Wiegmann, Alex. 2017. No Need for an Intention to Deceive? Challenging the Traditional Definition of Lying. *Philosophical Psychology*, 30(4), 438–457.
- Sorensen, Roy. 2007. Bald-Faced Lies! Lying Without the Intent to Deceive By. *Pacific Philosophical Quarterly*, 88(2), 251–264.
- Sorensen, Roy. 2010. Knowledge-Lies. *Analysis*, 70(4), 608–615.
- Staffel, Julia. 2011. Reply to Roy Sorensen, 'Knowledge-Lies'. *Analysis*, 71(2), 300–302.
- Staffel, Julia. 2016. Beliefs, Buses and Lotteries: Why Rational Belief Can't Be Stably High Credence. *Philosophical Studies*, 173(7), 1721–1734.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge University Press.
- Stokke, Andreas. 2013a. Lying and Asserting. *Journal of Philosophy*, 110(1), 33–60.
- Stokke, Andreas. 2013b. Lying, Deceiving, and Misleading. *Philosophy Compass*, 8(4), 348–359.
- Stokke, Andreas. 2018. Lying, Sincerity, and Quality. In: Meibauer, Jörg (ed), *The Oxford Handbook of Lying*. Oxford University Press.
- Sturgeon, Scott. 2008. Reason and the Grain of Belief. *Noûs*, 42(1), 139–165.
- Turri, Angelo, & Turri, John. 2015. The Truth About Lying. *Cognition*, 138, 161–168.

- Turri, John. 2011. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, **89**(1), 37–45.
- Turri, John. 2016a. *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Cambridge: Open Book Publishers.
- Turri, John. 2016b. Knowledge Judgments in ‘Gettier’ Cases. *Pages 337–348 of: Sytsma, Justin, & Buckwalter, Wesley (eds), A Companion To Experimental Philosophy*. New York: Wiley-Blackwell.
- Turri, John. 2017. Knowledge Attributions in Iterated Fake Barn Cases. *Analysis*, anx036.
- Unger, Peter K. 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- Weiner, Matthew. 2005. Must We Know What We Say? *Philosophical Review*, **114**(2), 227–251.
- Weiner, Matthew. 2007. Norms of Assertion. *Philosophy Compass*, **2**(2), 187–195.
- Wiegmann, Alex, Samland, Jana, & Waldmann, Michael R. 2016. Lying Despite Telling the Truth. *Cognition*, **150**, 37–42.
- Williams, Bernard. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton: New Jersey: Princeton University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Worsnip, Alex. 2017. Contextualism and Knowledge Norms. *In: Ichikawa, Jonathan (ed), The Routledge Handbook of Epistemic Contextualism*. Routledge.