

Lying, Belief, and Knowledge

Matthew A. Benton

1 Introduction

What is the relationship between lying, belief, and knowledge? One natural first pass idea is to understand lying in terms of knowledge: to modify a line from Mark Twain, “Lying is telling someone something you know ain’t so.”¹ But many will regard this natural idea to be incorrect from the start, insofar as they judge that one can lie even when one simply believes, falsely, the negation of what one asserted: on their widely-accepted view one can lie even if one mistakenly tells another the truth. In what follows I shall proceed as if this widely-held view should be respected as orthodox, though in §5 I shall revisit a version of the Twain-style knowledge account of lying to explore how plausible it might actually be.

Several philosophers debate the exact conceptual analysis of lying. Many accept an analysis with the traditional three-clause structure:

(TRAD) You lie to S just in case you:

- (i) You assert that p to S;
- (ii) You believe that p is false; and
- (iii) By asserting p , you intend to deceive S (in some specified way).

Much of the debate over such an analysis concerns the exact nature of the third condition,² in particular, whether the speaker intends for S to infer

¹See the quote about faith from the school boy, in Twain’s *Following the Equator*, ch. 12, “Pudd’nhead Wilson’s New Calendar” (1898).

²Different accounts specify the intent-to-deceive clause differently: see e.g. Mahon

something false about the proposition asserted, or about what the speaker herself believes, or perhaps that the former be inferred from an inference about the latter. However, many argue that a third such clause, requiring something about the speaker's intent to deceive, is not even necessary for an assertion to be a lie; for example, Sorensen (2007), Saul (2012), and Stokke (2013) defend definitions of lying which retain the first two clauses above, but abandons any such third condition. Even if such a third condition is not necessary, most commentators are primarily interested in lies which are intended to deceive, and one might approach the project of analyzing what a lie is by focusing on deceptive lies (cf. Fallis 2010, 4ff.). While such lies are worthy of such interest, we shall see that an account of lying is available which can explain with other resources the deceptive intent typical of lies.³

The main focus of this chapter is the relationship, if any, between knowledge and belief, which figure in the epistemic norm of assertion, and the moral/social norms against lying. I develop a broad account of the injunction against lying according to which, roughly, lying involves intentionally misrepresenting oneself at the level of what one believes or knows, which in turn represents one's epistemic relationship to the proposition asserted. Two such accounts will be articulated, one in terms of belief, and one in terms of knowledge. Approaching the project in this way enables us to see how a norm against lying may be derived from the epistemic norm of assertion. And it opens up a new perspective on how to understand lying which mirrors the epistemic expectations which hearers have toward speakers who make assertions in testimonial exchanges.

2008, and Lackey 2013. See Fallis 2010 for an overview.

³For the purposes of this paper, I shall largely set aside concerns about irony, hyperbole, or otherwise inaccurate claims involving scalar considerations: for example, when someone claims that he just swam a mile, or that it is 9 o'clock even though they know it is 8:59. Whether these ought to be classified as lies is a difficult question (though see Marsili 2014); but one might think that if so, this gives motivation for including a very specific clause about intent to deceive.

2 The Norm of Assertion

First some terminological preliminaries. *Assertion* is the linguistic act of committing outright to the truth of a proposition: by one's utterance one declares or claims or states⁴ that something is so. As such, assertion is regarded as the default speech act for outright utterances made in the declarative or indicative mood. Speakers use assertions to make unhedged claims about the way things are, such as when one states "London is the capital of the U.K.," or "Joanna is in her office."⁵ Amongst speech act theorists (e.g. Searle 1979, 2–5), it is common to distinguish assertion from other types of speech acts (for example from interrogatives, imperatives, or performatives, among others) by noting, first, that the point of assertion is to represent the world as being a certain way, namely, the way represented by the proposition asserted; second, that assertion has a word-to-world direction of fit (the utterance is made true by conforming to the way the world is); and third, that assertion of a proposition expresses the psychological state of belief in that proposition. On this approach, assertion expresses belief even when the speaker asserts insincerely by not believing what she asserts.⁶

What is meant by an *epistemic norm* of assertion? Assertion's epistemic norm is a social linguistic rule that one satisfy some relationship to that proposition for properly asserting: typically, it specifies a necessary condition on (epistemically) proper assertion, such as that one must know it, or believe it, or have adequate evidence for believing it, or that it be true. Philosophers often think of belief as the mental correlate of assertion: it is the mental act of outright commitment to a proposition as true, which normally triggers dispositions to act on that belief. But beliefs, and assertions, can be false; whereas *knowledge*, though normally regarded as involving

⁴Some theorists distinguish statements from assertions (Meibauer 2014, 71, citing the German distinction between *Feststellungen* and *Behauptungen*). In English, however, there is little to distinguish them.

⁵For lengthy overviews, see Jary 2010 and Pagin 2015; for shorter overviews, see MacFarlane 2011 and Goldberg 2015, Ch. 1.

⁶Searle 1979, 4. This importantly updates Searle's original position (Searle 1969, 64ff.), on which sincerity was required in order even to make an assertion.

belief (one's belief can amount to knowledge), is distinguished from mere belief in that one can only know facts, that is, true propositions. Though one can hold a false belief, one can only know truths. Because of this, it is often said that knowledge is "factive."⁷ (Moreover, knowledge requires some normative condition(s), for example that one's belief be well-formed, the product of a safe method, or fit one's evidence, etc. Epistemologists disagree on the exact nature of these conditions, which typically fall under the label of "epistemic justification"; we shall not consider these issues here.)

Grice's Maxim of Quality offers a candidate epistemic norm of assertion: its supermaxim is "Try to make your contribution one that is true," and its two specific submaxims are

1. Do not say that which you believe to be false.
2. Do not say that for which you lack adequate evidence.⁸

Grice's first submaxim is sometimes called the *sincerity* condition, the second the *evidence* condition. Strictly speaking, Grice's Quality supermaxim does not require that one's assertion be true. But if we interpret it as being concerned with truth (and not merely *trying* to assert truths), we may understand its submaxims as corollaries the following of which will help one fulfil the supermaxim.

Timothy Williamson (2000: Ch. 11) and others⁹ have defended knowledge as the epistemic norm of assertion (KNA), such that this norm is constitutively related to the speech act of assertion:

(KNA) One must: assert that p only if one knows that p .

⁷For more on factivity, see Benton (draft, Ch. 1).

⁸Grice 1989, 27. Benton (2016) argues that Grice's maxim of Quality is best understood as the knowledge norm, "KNA" in what follows.

⁹See especially Turri (2010, 2011, 2014a), Benton (2011, 2016, draft, Ch. 3), and Buckwalter & Turri (2014), among many others. For an overview of this literature and its critics, see Benton (2014, §1), and Pagin (2015, §6.2).

KNA imposes a strict duty, forbidding the combination that: one asserts p and one does not know that p . Because knowing involves believing, KNA forbids insincere assertions; and because knowing is factive, KNA forbids false assertions. On KNA, one does something epistemically improper when one asserts without knowing what one asserts, even if one thereby does something appropriate in, for example, the *prudential* or the *moral* sense. Much of what follows invokes KNA, but very little turns on this; many of the ideas developed here generalize to most other rival accounts of the norm.

Lying is a category that most naturally applies to assertion, rather than to presupposition or conversational implicature, or to other speech acts, for example, to suggesting that a proposition is or might be true, or asking questions concerning whether a proposition is true.¹⁰ Thus a natural route into the connections between the communicative practice of assertion and the norms against lying is by considering the epistemic norm of assertion.

3 Norms and Evaluative Dimensions

We can gain some purchase on the relationship between KNA as the norm of assertion and lying by considering how KNA will affect conversational patterns, and how it will generate expectations of conversational participants who implicitly grasp the relationship of KNA to assertion. In particular, if KNA is correct, assertions will typically represent their speakers as knowing what they assert;¹¹ for by engaging in the practice of assertion

¹⁰Does one do something akin to lying if one *predicts* that p when one knows that not- p ? This will depend on if one treats the force of predicting as on a par with asserting (e.g. for Stokke 2013, who thinks of assertion along Stalnaker's lines of proposing that a proposition be added to the conversation's common ground, predicting might well have the same force as asserting). For an epistemic norm of the speech act of prediction in terms of proper expectation, see Benton and Turri (2014, esp. 1862ff.).

¹¹Moore (1962, 277) and Unger (1975, 251ff.) endorse the thesis that asserting that p represents its speaker as knowing that p . The case for that thesis turns on conversational patterns and the kinds of criticism available, which any norm of assertion ought to

whose norm KNA its speakers implicitly grasp, a speaker would represent herself as conforming to KNA. Even if we allow that asserting instead represents something weaker about oneself—perhaps it only represents oneself as believing, or as having some adequate evidence for, what one asserts—then still, asserting what one knows one does not believe will involve intentionally misrepresenting oneself in some way.¹²

Does lying always involve misrepresenting what one believes? One gloss might be this:

(LIE) A lies to B if and only if:

- (i) A asserts that p to B,
- (ii) A believes that p is false (and doesn't believe, under the same guise, that p is true), and
- (iii) A intends by so asserting to represent herself as believing that p is true.¹³

The parenthetical clause of condition (ii) is needed to rule out as not being lies cases where one believes a proposition and its negation; this possibility, though typically overlooked, can occur (on a certain view of structured

well-explain. Williamson notes that the KNA norm subsumes the assertions-represent-knowing thesis under more general principles: “In doing anything for which authority is required (for example, issuing orders), one represents oneself as having the authority to do it. To have the (epistemic) authority to assert p is to know p ” (Williamson 2000, 252 fn. 6). Because the norm itself explains why asserters do represent this of themselves, the case for the representation thesis is also the case for the norm. Pagin (forthcoming, §6 fn. 20) disputes this by claiming that even here, “authority” has a non-normative reading, namely “*satisfying of sufficient conditions*, which need not be related to any norms”; but sufficient for what? Presumably an answer must be given in normative terms: that is, sufficient for *acceptable* or *appropriate* action undertaken as an authority.

¹²This is because people do typically believe what their evidence adequately supports; and plausibly in virtue of this, speakers express belief, or represent themselves as believing, when they assert. This creates a problem for Lackey's (2007) alternative Reasonable-to-Believe Norm of Assertion (which does not even require belief for proper assertion), particularly because she also endorses a 'Not Misleading Norm of Assertion' (2007, 615–617). For criticism, see Benton (2016, esp. fn. 27) and Turri (2014b).

¹³This gloss bears some resemblance to one considered by Davidson (1985, 88).

propositions) when one believes a proposition and its negation under different guises. Suppose I see the same object through two different windows, but don't realize that it's the same object. Through the right window, it looks gray, and through the left window it looks green. I come to believe both: *that object [seen through the right window] is gray*, and also *that object [seen through the left window] is not gray*. But if the object is in fact gray, and I tell someone (thinking of what I see through the left window) that the object is not gray, I haven't lied. But without the parenthetical clarification in (ii), LIE would count this as lying.

More generally, it might be possible for someone to know a proposition but, perhaps by misremembering or by habit or by some irrational process which they do not recognize, they come to believe its negation. (My wife accidentally leaves her mobile phone behind, and I find it; it looks much like my own. Later I need to call her at a moment when I happen to be holding her phone, and out of habit, using her own mobile phone which looks like mine, I begin to try to call her mobile number. For that moment, I plausibly believe that she has her phone, even though I know that she does not have it. If during that moment I tell someone that she has her phone, I am not lying.)

Some may suppose that LIE's condition (iii) is superfluous because they think of assertion as an intentional speech act requiring a speaker to intend to represent herself as being sincere, that is, as believing what she asserts. This way of thinking about assertion might well be correct; but it seems committed to the idea that one cannot ironically or metaphorically or playfully *assert* a falsehood so as to mean or convey instead its negation. On this view of assertion, irony involves cases where one *utters* that p , and "makes as if to"¹⁴ assert that p , but one is in fact asserting (roughly) that $\neg p$. Those who understand assertion in this way may leave off condition (iii), and insist that with ironic utterances a speaker does not lie, for she does not even fulfill condition (i).¹⁵ But for those who prefer a thinner conception of assertion according to which one can assert something

¹⁴See Grice (1989, 34, 53–54).

¹⁵See Stokke (2013, and "Lying, Sincerity, and Quality," this volume) for helpful discussion.

without intending to communicate it, or without intending to represent oneself as believing it, (iii) will be needed. On the thinner picture of assertion, (iii) can itself provide an interpretation of why standard cases of irony are not lies: for the hallmark of irony is that it typically doesn't involve the speaker intending to represent herself as believing the falsehood uttered (*mutatis mutandis* for cases of metaphor, hyperbole, and so on).

Condition (iii) need not involve the intent to deceive B concerning either p or what A believes, because A might know that B knows what A believes. So this gloss properly counts many of Sorensen's bald-faced lies and "knowledge-lies"¹⁶ as lies. For bald-faced lies and knowledge-lies are plausibly speech acts which are nevertheless used to represent oneself as believing what one asserts, even if everyone involved knows or thinks it likely that one does not believe it.¹⁷ Clause (iii) also enables a satisfying pragmatic account of why, typically, lies do involve the intent to deceive: most scenarios in which one wants to misrepresent oneself are also occasions on which one does so for deceitful purposes. On those occasions, (i)–(iii) obtain because one endeavours to deceive someone into thinking that one believes p , and typically as a means to getting them to believe that p .¹⁸ But because what is *typical* of lying may not be required for its definition, I suspect that a definition of lying need not include a clause concerning the intent to deceive. So I will proceed as if LIE is plausible enough to begin with as a general characterization of lying.¹⁹

Recall that KNA forbids the combination that one asserts a proposi-

¹⁶See Sorensen (2007 and 2010; cf. also Staffel 2011). The above account will count such lies as knowledge-lies when they uttered with the intent to prevent the addressee from knowing that what is asserted is untrue, but without the intent to deceive them into believing it (see Sorensen 2010, 610).

¹⁷An example of a bald-faced lie: a student who is caught red-handed cheating on an exam nevertheless denies it to the headmaster in his office, because denying it results in what the student deems a more preferable punishment than would confessing to the cheating (for similar examples see Lackey 2013 and Carson 2006, 290).

¹⁸One might think that one can lie while satisfying LIE's (i) and (ii) by asserting p whilst only intending to represent herself as being in a good epistemic position with respect to p (that is, regardless of whether she believes p). I don't think that's a lie, but I do find something objectionable about it: I shall describe what is objectionable about it in §4, under the notions of *negligent* and *vicious* assertion.

¹⁹Fallis (2010, 13 fn. 43; and 2013) seems to agree, commenting on Davidson (1985).

tion but one does not know that proposition. That is, KNA is logically equivalent to: *One must not: (assert that p & not know that p)*. If KNA is the norm of assertion, then we can plausibly derive a secondary norm of assertion on which one *should* assert only if one reasonably takes oneself to know. This follows from a quite general schema connecting norms of a certain structure to the conditions under which one should do what the norm permits. The schema reads:

(1) If one must (ϕ only if C), then one should (ϕ only if one has evidence that C),

where the “transition from ‘must’ to ‘should’ represents the transition from what a rule forbids to what it provides a reason not to do” (Williamson 2000, 245). If we suppose (as seems plausible) that evidence is what makes belief reasonable,²⁰ and if we prefer talk of reasonable belief over merely one’s having evidence, we may substitute these in thus:

(1*) If one must (ϕ only if C), then one should (ϕ only if one reasonably believes that C)

Substituting “assert p ” for “ ϕ ” and “one knows p ” for “ C ” yields the proposed derivation from KNA, which we may call KNA*:

(KNA*) If one must (assert p only if one knows p), then one should (assert p only if one reasonably believes that one knows p)

²⁰Epistemologists think of evidence in a variety of ways (some think of it as consisting of beliefs, or experiences, or propositions, or some collection of these; Williamson thinks of it as knowledge: 2000, Ch. 11). But epistemologists tend to agree that however it is understood, evidence is the kind of thing (perhaps not the only thing) which somehow supports hypotheses (typically, by making them more probable). So they tend to agree that a belief that p is reasonable if p is supported by one’s evidence.

Given KNA, one may apply *modus ponens* to KNA* to derive the following norm (call it KNA**): one should assert a proposition only if one reasonably believes that one knows it.

Note that KNA**, derived in this way from KNA and KNA*, might seem to impose an extremely strong standard on what it takes to be in a position to assert reasonably.²¹ But KNA** is much stronger than what we need to generate a useful derivative norm. For given KNA, one who endorses it (tacitly or explicitly) will tend to accept the following:

(G) One must: refrain from asserting if one believes or accepts that one doesn't know that *p*.

Assuming normal cognitive agents who tend to believe or accept what their evidence suggests, (G) will serve as a guidance norm for one who aims to conform to KNA, because following (G) will make one less likely to violate KNA.²² Agents who accept KNA will violate (G) on pain of either incoherence or *akrasia*: it would be incoherent if they accept KNA, prohibiting assertion of one does not know, but they also accept (against G) that they may on this occasion assert that which they believe they don't know; it would be *akratic* if they accept KNA but due to weakness of will, do anyway what they know that (G) forbids. Either way, they would be choosing to assert a proposition which they believe or accept to be in violation of a norm which they endorse.

²¹Though it is one which Williamson (2000, 256) seems to accept; how strong it actually is will depend on how hard it is to reasonably believe that one knows when one in fact knows; for example, it might be that knowledge that one knows isn't that uncommon. Nevertheless, violating KNA** is presumably less important than violating KNA itself: one might develop virtuous cognitive habits such that one tends to assert only what one knows, even though one rarely forms a belief about whether one knows.

²²I am also assuming that for such normal cognitive agents, their practice of believing what their evidence suggests makes the propositions believed objectively more likely to be true. For strange agents who (perhaps for some domain) constantly believe falsehoods more often than they believe truth, (G) won't help them much.

4 Lying and Vicious Assertion

One may distinguish evaluative dimensions of responsibility built up from any norm with a certain structure. Consider any moral, prudential (etc.) norm enjoining one to undertake an action ϕ only in circumstances C . If one endorses that norm on ϕ -ing, then one may be said to ϕ *reasonably* when one believes one is in C ; one may be said to ϕ *negligently* when one does not even consider whether one is in C ; and one may be said to ϕ *viciously* when one believes that one is *not* in C . For example, if a moral norm enjoins one to act only if one causes no one harm, then: one reasonably acts when one believes that by so acting one will cause no harm; one acts negligently when one acts without consideration of whether one's action will cause harm; and one acts viciously when one acts while believing that one's action will cause harm. Or for a legal example, suppose there is a legal rule enjoining one to send someone to prison only if they are guilty of breaking a law. A sentencer would thus act reasonably by sending someone to prison whom she believes to be guilty; she would act negligently by sending someone to prison when she doesn't consider whether or not they are guilty; and she would act viciously by sending someone to prison whom she believes is not guilty.

Given (G), and again on the simplifying assumption that a speaker grasps and endorses KNA, we can also distinguish between evaluative dimensions of responsibility in the domain of assertion. These layers of responsibility turn on one's reflective stance toward one's epistemic position concerning the asserted proposition p :

- (E1) One *reasonably* asserts that p when one believes that one knows p (that is, when one believes one conforms to KNA).
- (E2) One *negligently* asserts that p when one doesn't consider whether one knows p .
- (E3) One *viciously* asserts that p when one believes that one does not know p .²³

²³These distinctions owe much to Unger (1975, 260–262) and Williamson (2000, 245

These evaluative terms (*negligent, vicious*) aim to label the distinctively epistemic responsibility at play in assertion. For example, (E2) aptly captures something akin to Frankfurt's (2005) notion of bullshitting, that is, asserting without regard to whether one's claims are true.²⁴ Likewise, the viciousness of (E3) is primarily epistemic; so-called "prosocial" lies, though undertaken with high moral or social aims, nevertheless exhibit epistemic vice.

According to (E1), one can reasonably assert even if one violates KNA (and thereby asserts improperly): by asserting when one believes that one knows, one is in some sense less subject to blame if one violates KNA, for one does something right in virtue of trying to conform to KNA.²⁵ According to (E2), one negligently asserts when one fails to evaluate one's epistemic position with respect to p but goes ahead and asserts p anyway; even if one in fact fulfills KNA, one has done something a bit irresponsible. Finally, given (E3), one viciously asserts when one also believes that one asserts in violation of KNA. In the case of vicious assertion where one asserts that p in order to express what one thinks that one does not know, because one knows that one believes the opposite, we plausibly have a lie. That is, where a speaker asserts that p in order to assert what she herself realizes that she does not believe (because she knows she instead believes that $\text{not-}p$)—and where she does not flag by any other means (such as by winking to her interlocutor) that something special is going on—she lies. Lying is then a special case of vicious assertion.²⁶

and 255–256). Some, particularly those who think that KNA is too strong, will likely regard (E3), for example, as much too strong. Note however that (E3) fits well with our intuitive reaction to the knowledge version of Moore's paradox, namely assertions of the schema " p but I don't know that p ." And (E2) fits well with our reaction to an assertion of " p but I haven't considered whether I know p ."

²⁴For more, see Andreas Stokke, "Bullshitting," this volume.

²⁵Compare Williamson's (2000, 256) distinction between reasonably asserting and permissibly asserting, and also DeRose's (2009, 93–95) distinction between "primary propriety" and "secondary propriety."

²⁶Unger (1975, 260–262) gestures at such an idea, utilizing the asserting-represents-knowing thesis. Marsili (2014) claims that *graded-belief* lies, wherein one asserts that p

5 Knowledge and Lying

Above we considered a subject who lies by asserting p while knowing that she believes that not- p . Because she knows that she does not believe p ,²⁷ she thus knows that she does not know p . By asserting that p , she misrepresents herself as believing p when she does not; and (E₃) classified this as a vicious assertion, which our gloss LIE above counts as a lie.

LIE can seem plausible as a general characterization of lying, in part because its third clause invokes the intent to misrepresent oneself even though that intent may not extend to deceiving others. But perhaps the Twain-inspired knowledge slogan—that lying is telling someone something you know ain't so—should be revisited, for it makes possible an attractive account of lying which parallels the knowledge norm of assertion:

LIE-K A lies to B if and only if:

- (i) A asserts that p to B,
- (ii) A knows that p is false (and doesn't believe, under the same guise, that p is true²⁸), and
- (iii) A intends by so asserting to represent herself as knowing, or at least believing, that p is true.²⁹

Like LIE from §3, this knowledge account of lying counts bald-faced lies and “knowledge-lies” as lies. LIE-K also makes good on the idea from §4, that lying is a special case of vicious assertion. In fact, it makes characterizing that special case even easier: one viciously asserts p when one asserts

but only believes that p is probably false, intuitively also count as lies. On the accounts developed here (LIE above, and LIE-K below), they are not full lies, but they are still special cases of vicious assertion: one asserts the negation of what one thinks is probable, in order to assert what one realizes one does not (outright) believe.

²⁷We're assuming that she realises that she does not believe both that p and that not- p .

²⁸For reasons delineated in §3 above, discussing LIE.

²⁹Turri (2016) considers a similar “known-false” gloss on lying, but his does not include a third clause. Recall however that the third clause might be dropped on a thicker conception of what assertion is.

it to express p when one also believes that one doesn't know that p . But when one knows that p is false it will typically be quite easy to form the belief that one doesn't know p . So the special case of viciously asserting occurs when one believes that one doesn't know what one asserts because one knows its negation. Lying is thus the anti-paradigm of conforming to KNA: asserting what you know follows KNA; intentionally asserting the negation of what you know is lying.³⁰

This mirroring relationship between the KNA upholds the intuitive idea that lying involves violating a conversational norm (Fallis 2012, 577). It also suggests an elegant account of why it would be that lying is possible for assertion but not for conversational implicatures or other non-explicit messaging, which themselves often depend on assertion.³¹ If lying is conceptually related to the knowledge norm of assertion in the way envisioned above, it explains why one cannot lie that p by conversationally implicating that p : because conversationally implicating that p does not represent oneself as knowing p , one does not undertake the same epistemic responsibility by conversationally implicating that p as when one asserts that p with full force. This is partly because conversational implicature and other implicit messages make possible deniability of what was conveyed, something which is not available with outright assertions. Claiming "I never said that p " when one has merely conversationally implicated that p can typically get one off the epistemic hook for having conveyed that p ; but clearly, when what one said was that p , one cannot (truthfully!) deny that one said it.

LIE-K has the virtue of being more simple by needing no intent to deceive clause, a clause invoked by many traditional accounts of lying (compare the third clause of TRAD schema with which this chapter began). And when combined with KNA, it enables an account of why typical lies are misleading as well as epistemically problematic, even without an intent to

³⁰Notice also that endorsing the derived KNA** from §3 itself would provide a strong norm against lying in the sense given by LIE-K: for conforming to KNA** will ensure that one does not meet the conditions of LIE-K.

³¹See Fricker 2012 and Hawthorne 2012 for discussion. Adler 1997 argues that false conversational implicatures are not lies. Meibauer (2014, 113ff.) demurs, though he also allows that false implicatures could be called deceptive rather than lies (2014, 135).

deceive clause. Typical lies are misleading because, given KNA, they represent their speakers as knowing, and so believing, something which their speakers know to be false. And they are epistemically problematic because they frustrate hearers' cognitive efforts at gaining knowledge from others by way of testimony: indeed, when lies are taken up by hearers as the truth, they generate false beliefs which their hearers take on trust from their interlocutors to be knowledge. When testimonial uptake leads one to think that one knows something which is in fact false, and which one took on trust from someone else who was in fact the opposite of trustworthy, it's not just that one's epistemic position is compromised; the social goodwill we have toward others, especially those with whom we have ongoing relationships, is strained.

Notice that preface paradox scenarios can generate a case judged to be a lie by a belief account, but which is not judged a lie by the LIE-K account.³² Think of a writer who has painstakingly researched and written a non-fiction book, in which she makes numerous claims. She accepts each claim as true, though she recognizes that given the book's length and her own fallibility, at least one of her claims is likely to be false; and in the preface she admits as much. She thus has asserted (by writing the whole book) a large series of claims, but has also claimed that (at least) one of them is false.³³ Similarly, suppose S believes each of a large set of propositions $p, q, r, \dots n$, but which is a set such that S does not believe their conjunction. That is, S believes each of these propositions individually, but does not believe or assent to their conjunction: for S knows that one of the set's conjuncts is likely to be false, though S knows not which. For these reasons, S believes that the conjunction (call it 'CON') is false. And we shall stipulate that one of the conjuncts in that set is in fact false. If S asserts the conjunction CON (" p and q and r and ... n ") to someone, S is asserting a proposition which S does not believe (though S believes each of its conjuncts); indeed, S believes that CON is false.

Let us suppose that the intuitive response to such a case is that S's conjunctive assertion of CON is not a lie. On many belief accounts of lying,

³²Thanks to John Hawthorne for suggesting this case.

³³Cf. Makinson 1965.

such as *LIE* from §3, the assertion of the conjunction *CON* will presumably, but counterintuitively, count it as a lie. This is because *S* believes that *CON* is false. Even traditional schemas for defining lying such as *TRAD* will either have to deem it (counterintuitively) as a lie, or will have to explain the case as one in which *S* does not, by asserting the conjunction, intend to deceive the hearer into believing the conjunction. Yet the latter option seems a tall order, for asserting the conjunction looks like an act which is undertaken to put forth that conjunction to the hearer, even though the speaker believes that conjunction to be false. *LIE-K* appears to have an advantage here: for *LIE-K* does not even count the assertion of *CON* as a lie. This is because *S* believes, but does not know, that *CON* is false. Now, *S*'s asserting *CON* would count as a vicious assertion because *S* regards herself as believing that she does not know *CON*. But it does not meet *LIE-K*, nor is *S*'s asserting *CON* one of our special (lying) cases of vicious assertion. This is because *S*'s assertion of *CON* is not undertaken in order to express what *S* knows that she doesn't believe (*S* would know this by knowing that she instead believes that *CON* is false). Rather, *S* undertook asserting *CON* merely by, and for the purpose of, asserting individually the conjuncts of *CON*, each of which *S* believes.

6 Objections

LIE-K will be opposed by those who insist that one can lie while (mistakenly) telling the truth.³⁴ These traditionalists always invoke belief rather than knowledge in the second clause of such accounts because they have the strong intuition that it is still a lie when one wrongly believes a falsehood, and tells someone its negation: if I really believe that the grocer is closed today, and I tell you it's open because I want to frustrate you by sending you to a closed shop (and you go and find it open and happily get on with your shopping), these traditionalists maintain that I have still told a lie. Some recent empirical studies suggest, however, that this intuition is wrong: results from Turri and Turri (2015) suggest that in fact,

³⁴Contrast Carson (2006) is rather exceptional in arguing that lies must be false.

most people require falsity for a telling to be a lie. Such findings are not the last word; indeed, the case developed for LIE-K above proceeded on other grounds. But consistent with such findings, the proponent of LIE-K can claim that in the above case, when I told you the grocer was open, I merely tried (and failed) to tell a lie. To know whether you've lied, the LIE-K proponent will suggest, we must not only look at the speaker's inner mental life, but also at the way the world is.

The LIE-K proponent can thus be concessive here to the traditionalist who prefers an account along the lines of (TRAD) or our (LIE) above: what traditionalists get right is that the attempted liar has done something wrong (and it might be that the traditionalist is conflating the judgment that the attempted liar has done something wrong, with the judgment that she has lied). In particular, the attempted liar has engaged in vicious assertion, for he tells someone something which he himself believes he does not know, because he knows that he doesn't believe it (instead, the would-be liar believes its negation). That's a vicious speech act the likes of which, when choosing to engage in it, involves choosing to misrepresent one's own state of mind to someone else.³⁵ But it need not follow from this that the attempted liar has in fact lied.

Given the LIE-K account, is it possible to lie to someone even when one is not engaged in vicious assertion? If our claim above is correct, that lying is a special case of vicious assertion, then no, it would seem not possible. Suppose that Bill and Ashley are talking about France, where Bill would like to visit. And suppose that (i) Ashley asserts that *Lyon is the capital of France* to Bill, (ii) Ashley knows that this is false (and doesn't believe that it is true), and (iii) Ashley intends by so asserting to represent herself as knowing that it is true. To be plausible that Ashley is not engaging in vicious assertion, it would have to be the case that Ashley does not believe that she knows *Lyon is the capital of France* to be false. Now,

³⁵Notice that Lackey's (2007, 2013, and her contribution to this volume) cases of "selfless assertions" are classified here as vicious assertions which are not in fact lies (on LIE-K) if they are true. I find this result satisfying, for selfless assertions are intuitively not epistemically appropriate assertions, even if they are in some respects admirable; for criticism of relying on such cases, see especially Turri (2015).

it might be possible for Ashley to assert this to Bill without even forming a belief about what she knows concerning France's capital. She might be asserting this to Bill simply because it was the first proposition about France which, though false, popped into her head, and Ashley was just not reflective enough to assess her epistemic position with respect to *Lyon in the capital of France*. In our taxonomy of evaluative dimensions (E1)–(E3), Ashley here is asserting negligently (E2) rather than viciously. Note however that (E1)–(E3) are only sufficient conditions for reasonable, negligent, and vicious assertion, respectively; as such, some negligence might itself be vicious, particularly if one is intentionally acting recklessly. If one interprets the evaluative notions in this way, Ashley's lie might be merely negligent given (E2), but viciously so; and so the lie is still a case of vicious assertion even though it may not be vicious in the sense given by (E3).

But even if we regard Ashley's assertion as negligent without being vicious, it is hard to see how Ashley could be unreflective about her epistemic position toward this proposition yet also (iii) intend by her assertion to represent herself as knowing that *Lyon is the capital of France*. For by intending through that assertion to represent herself as knowing this, it looks like she at least has a cognitive perspective on it which would involve enough reflection to see that she doesn't know it, and in fact knows its negation. More generally, I suspect that it will be hard to find cases in which a speaker does intend to represent herself as knowing something when she knows its negation, yet does not actually believe that she knows its negation. One reason it will be hard to find such cases is that coming to form an intention to represent one's epistemic state as being a certain way plausibly engages one's reflective capacities on whether one's epistemic state is that way. This is plausible because cases of lying involve doing the opposite of what we normally do. Most of us assert knowledgeably, and thus truthfully, most of the time, and even when we do not, our default aim is truth rather than falsehood. Cases in which we opt against this habit and tell falsehoods require us to note mentally that we are telling some one something against what we believe, and such mental judgments will involve reflection both on what we think is true and why not to tell

someone the truth.³⁶

7 Conclusion

In this chapter I have considered the relationship which belief and knowledge bear to our understanding of lying. That relationship turns on the epistemic norm of assertion, and in particular, on the way in which belief and knowledge are required for proper assertion. Most of the results considered may be applied whatever is the correct epistemic norm of assertion, though the knowledge norm, KNA, has been widely defended. As such, I have used KNA to provide the background for articulating the normative epistemic dimensions to which we hold one another responsible in conversation. Epistemically reasonable, epistemically negligent, and epistemically vicious assertion provide a taxonomy for understanding the epistemic and social aspects of the wrongness of lying, as well as the wrongs of other speech which are not lying, but which are nonetheless misleadingly deceptive. The norms which may be derived from KNA offer a new perspective on how to understand lying: lying is knowingly doing, because one knows the opposite, that which the norm of assertion forbids. By intending for one's assertion to represent oneself as knowing, or at least believing, the opposite of what one knows to be true, one lies. Lying is the inverse of an assertion which is beyond epistemic reproach; lying is as bad as knowingly asserting is good.³⁷

³⁶Lying engages more areas of the brain, including areas involved with suppressing the truth, and with anxiety. See Langleben et al. 2002 and Langleben et al. 2005.

³⁷Thanks to Max Baker-Hytch, Rachel Fraser, John Hawthorne, John Turri, and Andreas Stokke for discussion and helpful comments. I am also grateful to two anonymous referees for useful feedback. This paper was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the John Templeton Foundation.

References

- Adler, Jonathan. 1997. "Lying, Deceiving, or Falsely Implicating." *Journal of Philosophy* 94: 435–452.
- Benton, Matthew A. 2011. "Two More for the Knowledge Account of Assertion." *Analysis* 71: 684–687.
- . 2014. "Knowledge Norms." *Internet Encyclopedia of Philosophy*, ISSN 2161–0002, <http://www.iep.utm.edu/kn-norms/>.
- . 2016. "Gricean Quality." *Noûs* 50: 689–703.
- . draft. *Knowledge and Language*.
- Benton, Matthew A. and Turri, John. 2014. "Iffy Predictions and Proper Expectations." *Synthese* 191: 1857–1866.
- Buckwalter, Wesley and Turri, John. 2014. "Telling, Showing, and Knowing: A Unified Theory of Pedagogical Norms." *Analysis* 74: 16–20.
- Carson, Thomas L. 2006. "The Definition of Lying." *Noûs* 40: 284–306.
- Davidson, Donald. 1985. "Deception and Division." In Jon Elster (ed.), *The Multiple Self*, 79–92. Cambridge University Press.
- DeRose, Keith. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.
- Fallis, Don. 2010. "Lying and Deception." *Philosophers' Imprint* 20: 1–22.
- . 2012. "Lying as a Violation of Grice's First Maxim of Quality." *Dialectica* 66: 563–581.
- . 2013. "Davidson was Almost Right about Lying." *Australasian Journal of Philosophy* 91: 337–353.
- Frankfurt, Harry. 2005. *On Bullshit*. Princeton: Princeton University Press.
- Fricker, Elizabeth. 2012. "Stating and Insinuating." *Proceedings of the Aristotelian Society, Supplementary Volumes* 86: 61–94.

- Goldberg, Sanford C. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.
- Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Hawthorne, John. 2012. "Some Comments on Fricker's 'Stating and Insinuating'." *Proceedings of the Aristotelian Society, Supplementary Volumes* 86: 95–108.
- Jary, Mark. 2010. *Assertion*. Palgrave Studies in Pragmatics, Language, and Cognition. Basingstoke: Palgrave Macmillan.
- Lackey, Jennifer. 2007. "Norms of Assertion." *Noûs* 41: 594–626.
- . 2013. "Lies and Deception: An Unhappy Divorce." *Analysis* 73: 236–248.
- Langleben, Daniel D., Loughhead, James W., Bilker, Warren B., Ruparel, Kosha, Childress, Anna Rose, Busch, Samantha I., and Gur, Ruben C. 2005. "Telling Truth from Lie in Individual Subjects with Fast Event-related fMRI." *Human Brain Mapping* 26: 262–272.
- Langleben, D.D., Schroeder, L., Maldjian, J.A., Gur, R.C., McDonald, S., Ragland, J.D., O'Brien, C.P., and Childress, A.R. 2002. "Brain Activity during Simulated Deception: An Event-Related Functional Magnetic Resonance Study." *NeuroImage* 15: 727–732.
- MacFarlane, John. 2011. "What is Assertion?" In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 79–96. Oxford: Oxford University Press.
- Mahon, James. 2008. "The Definition of Lying and Deception." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/lying-definition/>.
- Makinson, D. C. 1965. "The Paradox of the Preface." *Analysis* 25: 205–207.
- Marsili, Neri. 2014. "Lying as a Scalar Phenomenon: Insincerity along the Certainty-uncertainty Continuum." In Sibilla Cantarini, Werner Abraham, and Elisabeth Leiss (eds.), *Certainty-uncertainty – and the Attitudinal Space in Between*, 153–173. Amsterdam/Philadelphia: John Benjamins Publishing.

- Meibauer, Jörg. 2014. *Lying at the Semantics-Pragmatics Interface*. Berlin: De Gruyter Mouton.
- Moore, G.E. 1962. *Commonplace Book: 1919–1953*. London: George Allen & Unwin.
- Pagin, Peter. 2015. "Assertion." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Spring 2015 edition.
- . forthcoming. "Problems with Norms of Assertion." *Philosophy and Phenomenological Research*.
- Saul, Jennifer Mather. 2012. *Lying, Misleading, and What is Said*. Oxford: Oxford University Press.
- Searle, John R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- . 1979. *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.
- Sorensen, Roy. 2007. "Bald-Faced Lies! Lying Without the Intent to Deceive." *Pacific Philosophical Quarterly* 88: 251–264.
- Sorenson, Roy. 2010. "Knowledge-lies." *Analysis* 70: 608–615.
- Staffel, Julia. 2011. "Reply to Roy Sorenson, 'Knowledge-lies'." *Analysis* 71: 300–302.
- Stokke, Andreas. 2013. "Lying and Asserting." *Journal of Philosophy* 110: 33–60.
- Turri, Angelo and Turri, John. 2015. "The Truth about Lying." *Cognition* 138: 161–168.
- Turri, John. 2010. "Prompting Challenges." *Analysis* 70: 456–462.
- . 2011. "The Express Knowledge Account of Assertion." *Australasian Journal of Philosophy* 89: 37–45.
- . 2014a. "Knowledge and Suberogatory Assertion." *Philosophical Studies* 167: 557–567.

———. 2014b. “You Gotta Believe.” In Clayton Littlejohn and John Turri (eds.), *Epistemic Norms: New Essays on Action, Belief, and Assertion*, 193–200. Oxford: Oxford University Press.

———. 2015. “Selfless Assertions: Some Empirical Evidence.” *Synthese* 192: 1221–1233.

———. 2016. *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Cambridge: Open Book Publishers.

Unger, Peter. 1975. *Ignorance: A Defense of Skepticism*. Oxford: Clarendon Press.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.