

Knowledge is the Norm of Assertion

Matthew A. Benton

WHEN WE COMMUNICATE with language, we represent ourselves in several ways. Various speech acts carry with them standard ways in which speakers thereby represent themselves. For example, imperatives, such as *commanding* someone to do something, represent the speaker as having, and exercising, some sort of authority over them. *Requesting* that someone pass the salt represents the speaker as wanting the salt, or at least wanting to be passed it, and perhaps also that the addressee is near enough to do so. Similarly, interrogatives, such as *asking* about a friend's whereabouts, typically represent the speaker as not already knowing, or even believing, an answer.

Assertion is likewise a type of speech act, one typically engaged in by uttering an unqualified declarative sentence. In asserting, one represents something about oneself as well as about the world: not only does asserting put forth a claim as true, but also, in asserting, a speaker thereby represents themselves as *knowing* the claim they've asserted. The stability of the linguistic data which support this representational idea has been coupled with intuitive normative judgments about particular cases. Many philosophers have suggested that the best explanation of this data is that the speech act of assertion is tightly connected to a norm or rule on such acts, according to which, roughly: one should not assert that p unless one knows that p . The present chapter summarizes the case made so far, and argues further that the speech act of assertion has knowledge as its central norm.¹

In section 1, I situate the knowledge norm, and provide some rival accounts. Section 2 provides an overview of the data which any such norm aims to explain. Section 3 considers some objections to the knowledge norm,

¹See Williamson 2000, Chap. 11, and Turri 2016b for the most thorough treatments.

and some complaints which masquerade as objections. Finally, in section 4, I offer some concluding thoughts.

1 Formulating the Knowledge Norm

Most agree that there is a norm with epistemic content which governs assertions, the default speech act made using unqualified declarative sentences, for example, “Houston averages more annual rainfall than Seattle.” Though the basic idea was embraced by many previous philosophers,² contemporary formulations of the knowledge norm take the form of a necessary condition³ on proper assertion, where the dimension of propriety tracks the epistemic position of a speaker to make the assertion. The most succinct version of the norm simply says that you should know what you assert. A more formal version is given by Williamson (2000, 241):

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

KNA thus forbids asserting when one does not know the proposition one asserts; equivalently, it permits asserting only when one knows what one asserts. As Williamson thinks of it, KNA should be understood “as giving the condition on which a speaker has the *authority* to make an assertion. Thus asserting p without knowing p is doing something without having the authority to do it, like giving someone a command without having the authority to do so” (2000, 257).⁴ As such, one should not assert p without having the requisite

²Moore 1962, 227 and 277; Austin 1961, 45; Unger 1975, Chap. 6; Slote 1979.

³Some also discuss whether knowledge is *sufficient* for (epistemically) proper assertion: see Lackey 2011 and 2016, and Benton 2016a for more.

⁴It is a further question what the status of this norm is, or how it fits into a fuller account of assertion. Williamson 1996/ 2000, Chap. 11, argues for a “simple” account (2000, 241–242), on which the norm is *constitutive* of the speech act of assertion, on analogy with the way that rules constitute a game, where grasp of the proper move within the rules implies mastery of the applicability of that rule (see Maitra 2011 for criticism). Williamson also claims the norm is individuating of assertion, thus unique to it (see Turri 2014, and Simion and Kelp 2020 for more). It should be noted, however, that several proponents of a knowledge norm on assertion do not feel the need to endorse some of these claims (e.g., Turri 2018 argues for it as a sustaining rule; Sosa 2010 and 2015, 170–171, as a teleological performance norm;

authority to assert p , namely, knowing p . The idea is that assertions are always governed by KNA,⁵ which is why asserting represents one as knowing.⁶

Rival accounts of the norm typically appeal to other notions of epistemic interest, such as a justification or credibility or supportive reasons norm (Lackey 2007, Douven 2006, McKinnon 2013), a justified belief norm (Kvanvig 2009), or a truth norm (Weiner 2005), among others. For example, Weiner’s truth norm says that one should assert that p only if p is true (and that this would provide a derivative norm requiring one to have evidence for its truth). Or, Lackey’s alternative norm says that one should assert that p only if it is reasonable (given one’s evidence) for one to believe that p , even if one doesn’t so believe. In what follows I focus mainly on the arguments presented in favor of KNA; some objections are considered in section 3.

2 Approaching the Data

The case for KNA derives from multiple strands of data which KNA explains well; indeed, it seems to offer the best unified explanation of those strands. These strands are often sorted into data from problematic conjunctions, several interrelated conversational patterns, and intuitive judgments of permissibility, excusability, and blame. Additional evidence comes from structurally comparable data from ascriptions of knowing and showing how.

or in functionalist terms, e.g., Turri 2016c, Kelp 2018). And some alternative approaches to individuating assertion can arguably nevertheless accept that knowledge is the central norm of assertion (see MacFarlane 2011: several of the views which he contrasts with Williamson’s constitutive account can similarly endorse the knowledge norm).

⁵They will also be subject to other norms of prudence or morality in a given context.

⁶Some philosophers, such as Williamson (2000, esp. 252, n. 6, and 257), eschew an account in terms of representation in favor of a more general account in terms of acting on a kind of authority (in this case, epistemic authority); whereas others, like van Elswyk 2021, defend the knowledge representation effect for declaratives in semantic terms (what he calls PARENTHETICALISM), and thus he needn’t invoke a knowledge norm, which nevertheless could be coupled with PARENTHETICALISM). For norm views with and without representational accounts, see Benton and van Elswyk 2020, 250–253.

2.1 Faulty conjunctions

Initial discussions of a norm of assertion standardly cite the strangeness of asserting Moorean paradoxical conjunctions,⁷ particularly the belief and knowledge versions:

(1) # It is raining but I do not believe it is.

(2) # It is raining but I do not know it is.

Notice that commuting the conjuncts of such assertions (or those discussed below) does not make them sound any better. The infelicity of such assertions (which I mark with “#”) isn’t merely one of surprise given that they rarely are made in normal conversation.⁸ It is rather that they positively clash: upon hearing them, it is unclear what to believe from such a speaker: should one trust the sincere asserter of (1) or (2) that it is raining?

What Moore and others found interesting about such conjunctions is that they might well be true; and thus any infelicity arising from asserting them cannot be due to semantic inconsistency, in the way that asserting a contradiction would be, or even a conjunction of a claim and something entailed by it (for example, “This is a square but not a shape”). Further evidence that semantic inconsistency is not in play comes from our ability to entertain the truth of each conjunct, the ease of embedding (1) or (2) within the antecedent of a conditional.⁹ So diagnosing their infelicity seems to depend on a certain pragmatic effect of making an assertion. Explaining it using the representational language from earlier: first, asserting that p represents oneself as knowing that p . And so by asserting the conjuncts of (2), one represents oneself as knowing that it is raining, and as knowing that one does not know

⁷Perhaps more aptly named after MacIver 1938, who discussed them first.

⁸After all, we are constantly hearing novel sentences or amusing ones which we can understand or at least entertain: the surprise induced by “Beyoncé is a vegetarian chipmunk” is not due to incoherence. (Thanks to Jason Stanley for this example.)

⁹That is, “If it is raining but I do not know it, then I will get soaked when I go out” is an understandable and a felicitous conditional. So it isn’t that the utterance of such conjunctions in all linguistic constructions sounds bad, but rather those where the conjunction is asserted outright.

this. But one cannot know that it is raining and also know that one does not know that it is raining (because knowledge is factive). So the infelicity arises from what one represents about oneself by asserting the first conjunct, and what one explicitly disavows as being so in the second conjunct. And the first claim, that asserting represents one as knowing, would be explained by KNA: for if there is a norm of permission on assertion requiring knowledge, in asserting one thereby represents oneself as satisfying that requirement. Finally, because knowing plausibly involves believing, a parallel argument can explain the infelicity of (1) in terms of KNA.

Yet more faulty conjunctions are also best explained by KNA. For example:

(3) # I hope that John is in his office, but he is not there.

(4) # I regret that they lost, but I don't know whether they lost.

(3) contains no knowledge disavowal, whereas (4) contains no conjunct asserting the proposition at issue. Yet with some plausible supplementation, KNA can also explain them. Hoping that p is somehow incompatible with knowing whether p .¹⁰ But then, given KNA, conjoining a self-ascription of hope in a proposition will conflict with an outright assertion of that proposition or its negation, because the assertion represents one as knowing it. Similarly, in (4), asserting with the factive predicate *I regret that* implies, or has as its precondition, that one knows the proposition regretted. Thus asserting its first conjunct represents one as knowing it; but as with (2), this conflicts with disavowing knowledge of in the second conjunct.

Noteworthy here is that other candidate norms on assertion which invoke lesser epistemic requirements such as justification, or justified belief, or truth, do not as easily explain what is wrong with the above conjunctions. For it seems to most philosophers perfectly plausible that I could, with regard to (1), say, be justified in believing that it's raining while also justified in believing that I don't know this.¹¹ Or with regard to (3), I might be justified in believing

¹⁰See e.g. Benton 2021.

¹¹Though not if knowledge is also the norm of belief; if so, then both believing and asserting (1) or (2) can be diagnosed as faulty in terms of knowledge.

that I hope John's in his office, but also justified in believing he's not there. So any such alternatives to KNA will have to appeal to other resources to explain why such assertions seem to come off so poorly, or why they seem to lack much communicative value.

2.2 Conversational patterns

More thoroughgoing data best explained by KNA comes from conversational patterns, which directly or indirectly invoke the standard of knowledge. These are often presented separately from the paradoxical conjunctions considered above, but as we shall see, some of these data are related, and explainable only by KNA.

*Prompts.*¹² To prompt an assertion, one may ask, "What time is the meeting?" or "Do you know what time the meeting is?" Both seem to be used interchangeably by speakers, and each is understood by hearers to be practically equivalent: each is asking for the meeting time. Why would this be? KNA can explain this by noting that knowledge is the standard for permissibly asserting in answer to such questions; so the former question requests an assertion which, given KNA, implicitly expects a knowledgeable answer. It is practically interchangeable with the latter question, since the latter directly asks for a knowledgeable answer, by citing that standard.¹³

Abstentions. When prompted for an assertion, one can decline to reply by citing one's lack of knowledge: "I don't know." And one can similarly abstain by saying, "I can't say," or "I can't tell." Indeed, one could disavow such knowledge while also indicating that one can't assert. The deontic modal "can" here flags whether one has permission so to assert. KNA easily explains all this: for in citing one's lack of knowledge, one explains why one cannot respond with an outright assertion; and claiming that one lacks the permission so to respond implies that it is because one lacks the relevant epistemic standing to do so, namely, because one does not know the answer.

¹²For these first three patterns, see especially Turri 2010 and 2016b.

¹³By contrast, we do not often prompt with questions like "What do you (reasonably) believe about...?" When we do use different prompts, like "Do you have any idea what...?", they signal an invitation for something weaker than outright assertion, such as a guess or hedged assertion. See below for hedges.

This explanation is reinforced by further faulty conjunctions involving appeals to such permission. For example,

(5) # Your case isn't still being reviewed, but I can tell you that it is.

(6) # I don't know whether your case is still being reviewed, but I can tell you that it is. (Turri 2011, 39)

KNA can offer a comparable explanation of such conjunctions to that given in section 2.1 above. By asserting a proposition you represent yourself as knowing it, but for (5) this seems to conflict, given KNA, with then claiming that you can tell them its negation. Similarly, for (6), by denying knowledge of a certain proposition you represent yourself, given KNA, as not being able to assert it; but this conflicts with the second claim that you can do so.

Challenges.¹⁴ One can typically, and appropriately, respond to an assertion with the question, "How do you know?" Yet asking this is normally regarded perfectly acceptable and relevant to a conversation, even though the speaker may not have explicitly referenced anyone's knowledge in their assertion. And often, such a question can be used as an implicit challenge to the authority of the assertor. Why would this question then be relevant, and enjoy a default propriety in most conversations; and why can it be used to implicitly challenge their authority? Given KNA, one's assertion is only proper if one knows what one asserted; so it will typically be relevant to ask how one knows.

Moreover, one can more more aggressively challenge an assertor by asking "Do you (really) know that?" More aggressive still is, "You don't know that!" In both cases, the speaker is put on the spot to defend their epistemic position with respect to their assertion, and if not well defended, the conversational participants will likely proceed as if the conversational common ground no longer accepts the proposition asserted. What explains all this? KNA can explain the increasing aggressiveness of such questions as follows: if "How do you know?" implicitly challenges a speaker's authority, "Do you know that?" *explicitly* challenges it; whereas "You don't know that!" *explicitly rejects* their

¹⁴See Austin 1961, 45; Unger 1975, 263–265; Williamson 2000, 252–253.

authority. Successful substantiation of one's knowledge is the clear standard for whether the conversation proceeds as if the asserted claim stands. (This is not to equate substantiating that one knows with actually knowing or having that epistemic authority; and for conversational purposes, defending one's claim can sometimes be difficult. What is noteworthy is that interlocutors will make a judgment about one's knowledge, and if deemed not knowing, the conversational "score" normally won't reflect one's asserted claim.)

It is often overlooked how the data from challenge questions relate to the faulty conjunctions discussed earlier. Yet the evidence from Moorean-paradoxical constructions like (2) and the evidence from challenge questions are related in a way that only KNA is well-positioned to explain. For the "How do you know?" challenge can elicit a *de facto* Moorean paradox within a conversational context:

A: It is raining.

B: How do you know?

A: Oh, I don't.

B: What?

A: Still, it's raining.

B's question effectively puts A into a potential Moorean predicament. So it looks like any explanation of what is problematic about the Moorean conjunction ought also to explain why the challenge questions are so apt, and vice versa: and even better, they ought to be given the *same* explanation. KNA does just that: because assertions represent their speakers as knowing, A's assertion invites the supposition that A knows; likewise, because assertions represent their speakers as knowing, any flat-out asserted conjunct of a Moorean sentence invites the supposition that its asserter knows it. Thus KNA offers a unified explanation of both data.¹⁵

Hedges. One strand of data that appears best explained by KNA concerns hedging behavior and our normative interpretations of them.¹⁶

¹⁵This is significant because KNA's competitors fare poorly at providing a unified account of the challenge questions and the Moorean conjunctions; indeed, they handle the data in a fragmented way, citing distinct considerations for each.

¹⁶Some of this material is covered more in Benton and van Elswyk 2020.

Speakers can append attitudinal expressions like *I think*, *I believe*, or even *I hope* to a declarative in order to convey a weakened epistemic position, thus:

- (7) Amia went home early.
- (a) Amia went home early, I think.
- (b) Amia went home early, I believe.

That is, in asserting (7) one represents oneself as being in a stronger epistemic position than one who utters (7a) or (7b). Likewise, adverbials like *reportedly* or *evidently*, or modals such as *may(be)*, *perhaps*, *might*, or *it's possible* can be used to similar effect.¹⁷ For example,

- (c) Amia reportedly went home early.
- (d) Perhaps Amia went home early.

It seems clear that each of these hedges serve to weaken the commitment from whatever the norm requires of outright assertion; given KNA, hedging distances the speaker from knowing the complement proposition which is the claim under discussion.

That knowledge is the norm is supported by two features of such hedging behavior. First, each of (a)–(d) above permit adding a conjunct disavowing knowledge of the proposition that Amia went home early. For example, such additions to (b), (c), and (d) yield the acceptable conjunctions (note well the aptness of *but* to conjoin):

- (b¬K) Amia went home early, I believe, but I don't know that she did.
- (c¬K) Amia went home early, reportedly, but I don't know that she did.
- (d¬K) Perhaps Amia went home early, but I don't know.

Yet by comparison, if one were instead to add a conjunct claiming knowledge (even when conjoining with *indeed*), the assertion will come off as oddly

¹⁷See especially Hawthorne 2004, 23–31 in his defense of KNA.

problematic. For if one claims to know in the second conjunct, it seems bizarre to have hedged in the first conjunct:

(b&K) ? Amia went home early, I believe; indeed, I know that she did.

(c&K) ? Amia went home early, reportedly; indeed, I know that she did.

(d&K) ? Perhaps Amia went home early; indeed, I know this.

Similar data can be generated using other hedged constructions, of course. This kind of evidence supports KNA in at least two ways. On the one hand, the hedged claims always seem compatible with disavowing knowledge, where both the hedged conjunct and the knowledge disavowal serve to explain why the speaker didn't simply unqualifiedly assert: each implicitly suggests that knowledge is what would've been needed in order to unqualifiedly assert it. Yet on the other hand, hedging feels out of place when one also claims knowledge, which is to be expected if one's having satisfied the norm of assertion absolves one from the need to hedge.

Parentheticals. A related kind of evidence for KNA involves a pattern found by using attitudinal expressions in parenthetical position as in (a) and (b). In each case, *I think* or *I believe* can take a fronting main clause position, or parenthetical position, including sentence-final:

(a) Amia went home early, I think.

(b) Amia, I believe, went home early.

But *I know*, though it can be used felicitously in main clause position ("I know that Amia went home early"), sounds odd and overly redundant in parenthetical position:

(k) ? Amia, I know, went home early.

(k) ? Amia went home early, I know.

So the sorts of attitudinal expressions which uniformly allow one to hedge against the primary proposition (that Amia went home early) are also those

which acceptably take on parenthetical position; whereas the attitude term specifying the KNA's content, *know*, sits redundantly in parenthetical position. (One might include contrasting information to make it non-redundant, such as adding to (k) above: "...but I don't know about the others." Of course, KNA can explain such exceptions as well.)

Furthermore, *know* also marks the difference between acceptable hedged claims conjoined with self-disavowals of knowledge, and redundantly strange conjunctions of hedged claims with self-attributions of knowledge. In other words, these hedging expressions cluster around the notion of knowledge and are applied rightly when distancing oneself from knowing, but applied wrongly when conveying or claiming knowledge for the speaker. The best explanation of these patterns is plausibly that knowledge sets the standard for proper assertion.¹⁸

Mutual reasoning. An overlooked argument from Adler 2009 shows that KNA explains cases where a hearer responds to a speaker's assertion with a claim that would *lower* the probability of the speaker's assertion, were its probability less than 1. Yet the hearer's response is not treated as undermining the speaker's claim:

... let the speaker assert that Miss Scarlett did the foul deed or the weapon was a wrench ($F \vee W$), expressive of his corresponding belief. The hearer responds by asserting that the wrench is too heavy for Miss Scarlett to have wielded it, $\neg(F \& W)$. In accord with the knowledge norm, neither speaker nor hearer treat $\neg(F \& W)$ as undermining evidence of $F \vee W$, but as a complementary contribution. (Adler 2009, 408)

B's reply ought to lower the probability of A's claim, for it removes one way in which A's (inclusive) disjunction could be true. But everyone will regard B's claim as complementary to A's rather than undermining of it, inviting the inference that only one of the disjuncts holds: either Miss Scarlett did the deed another way, or that another criminal used a wrench. Similar results apply in the simplest case of disjunctive syllogism, such as if B had instead replied that "Miss Scarlett did not do the foul deed." KNA can explain this speaker-hearer

¹⁸For more, see especially Benton 2011, Blauw 2012, and van Elswyk 2021.

harmony better than, say, a rational credibility norm: for given KNA, A's disjunctive claim conveys that A knows the disjunction but doesn't know which disjunct is true; and B's reply, rather than being understood as a challenge to A's disjunction, lowering its probability, is rightly interpreted as a supporting premise for mutual reasoning. B's contribution adds knowledge which refines the disjunction from inclusive to exclusive.

2.3 Permission, excuse, and blame

Some KNA theorists have appealed to intuitive judgments about the impermissibility of lottery assertions: to many, it seems improper to assert outright that you will lose a lottery for which you hold a ticket (before having heard the winner announced), even though your odds of winning make it exceedingly probable that you will lose.¹⁹ Many also find it plausible that one does not know that such a ticket will lose. KNA proponents explain the first point in terms of the second: the reason it is inappropriate for one to make such lottery assertions, absent special knowledge about the lottery being rigged, is that one does not know that the ticket will lose. But some people do not have such strong judgments about the impermissibility of lottery assertions, even those who otherwise defend KNA (e.g., Turri 2011, 37 n. 1). So such data seems less probative.

Yet there remains a set of judgments arising from taking the first-person perspective of deciding what and when to assert. When believing outright a proposition, many will feel this to be phenomenologically just like treating oneself as knowing; and thus the decision about whether to assert outright will be guided by consideration of whether one knows. This seems apt because any norm, KNA or otherwise, insofar as one aims to conform to it, will generate secondary norms of guidance. If KNA is correct, then the secondary norm derived from it will require that one refrain from asserting outright when one judges oneself as not knowing. (Note that this in turn provides an elegant explanation of the hedging patterns discussed above, and how these convey the responsibility of careful speakers.)

¹⁹Williamson 2000, 246–252; Hawthorne 2004, 21–23, among others. Note that some philosophers mistakenly think of lottery judgments as comprising all or a significant portion of the evidence for KNA (e.g. Lackey 2007, 611–613; Papineau 2021, 5329).

This structural feature of norms is fully general. For any norm of prohibition that requires one to do something only in conditions C, its secondary guidance norm will enjoin one to refrain from so doing whenever one regards oneself as not in C.²⁰ (Likewise for a different sort of norm that obliges one to do something in C, its secondary norm will enjoin one to do it whenever one regards oneself as in C.) Recognizing this distinction between layers of propriety enables the KNA theorist to explain our judgments about cases where one reasonably asserts without knowledge, where it seems one nevertheless has, in some sense, acted appropriately: “One may reasonably do something impermissible because one reasonably but falsely believes it to be permissible. In particular, one may reasonably assert p , even though one does not know p , because it is very probable on one’s evidence that one knows p ” (Williamson 2000, 256). Thus violations of KNA will be impermissible, but might nevertheless be reasonable, and thus need not be blameworthy. For assessments of blame typically track reasonableness, not permissibility.²¹ Indeed, when one does something wrong or impermissible, but they reasonably thought they were acting permissibly, we will be inclined to excuse them for the impermissible act, which also mitigates how blameworthy they are. For many such cases then, our willingness to blame or excuse can be explained in terms of secondary propriety or whether one acted reasonably given one’s grasp of the norm.

2.4 Knowing and showing

Finally, there is a striking parallel between the above data with respect to assertion, and related patterns with respect to pedagogical instances of showing how to do something (Buckwalter and Turri 2014; Turri 2016b, 21–25). In brief, just as knowledge is the norm of assertion (KNA), knowing how is also the norm of showing how. This is supported by comparable faulty conjunctions, and hedges:

²⁰Williamson 2000, 245 and 256 draws the distinction between “permissible” and “reasonable” assertion; DeRose 2009, 94–95 calls them “primary” versus “secondary” propriety.

²¹For more, see Kelp and Simion 2017 and Williamson *forthcoming*. For the dimensions of reasonableness, negligence, and viciousness (knowingly acting impermissibly), see Benton 2019, 127–128.

(8) # I don't know how to do it, but here is how it is done.

(9) # I can show you how to do it, but I don't know how.

These conjunctions sound quite flawed, in the way Moorean paradoxical conjunctions do. But inserting an appropriate hedge renders them acceptable:

(10) I don't know how to do it, but I think it's done something like this.²²

There is also parallel evidence from apt prompts, abstentions, challenges, and responses. By attempting to (or offering to) show one how to, say, ride a bicycle, one thereby represents oneself as knowing how to do so. The prompts "How do you do this?" and "Do you know how to do this?" as well as "Can you show me how to do this?" are all interpretable as either indirect or direct requests to be shown how. Not knowing how is sufficient excuse to abstain from fulfilling such requests. And one only properly responds affirmatively to such prompts if one indeed knows how and further, can show them how. There is a similar range of aggressiveness to the challenges, "How do you know how to do that?", "Do you (really) know how?", and "You don't know how to do that!" And so on.²³

Thus the above sections show the wide range of linguistic data and corresponding judgments which seem to be best explained by KNA. And such data are not concocted merely from armchair philosophizing: most of them are confirmed by experimental testing using thousands of subjects.²⁴ Not only does KNA provide elegant explanations of each such strand of data, from faulty conjunctions, conversational patterns, evaluative judgments, and more; it also offers a unified explanation of such data in terms of the knowledge norm, rather than needing to appeal to external resources to explain any portion of them.²⁵

²²Buckwalter and Turri 2014, 18.

²³Note that this datum is neutral with respect to whether knowledge-how is reducible to a sort of knowledge-that: see Buckwalter and Turri 2014, 19.

²⁴For experimental work from cognitive and social science testing this data, see Turri 2015, 2016b, 2016d, and 2017b.

²⁵Green 2017 and Haziza forthcoming argue that there is also a knowledge norm of im-

3 Some Objections and Complaints

Many philosophers have argued that their intuitive judgments about cases of asserting without knowing provide strong evidence against KNA. They judge that a speaker who asserts with only a justified belief that p speaks appropriately and permissibly, even if they assert p falsely or fail to know p (Douven 2006, 476–477; and Lackey 2007, 603, among others²⁶). In particular, they regard a speaker who justifiably believes that p , but unluckily given their evidence, p is false, who nevertheless asserts p , speaks blamelessly, for such a speaker is not properly criticizable for their assertion. But if blameless and not properly criticizable, such critics reason, a speaker cannot have violated the norm of assertion. So the norm of assertion cannot even require truth, let alone knowledge (which is factive, entailing truth). These judgments strike such philosophers as important objections to KNA, which drive them to defend a weaker norm, such as a justification or credibility norm, even though in so doing they are less able to explain all the data highlighted in section 2.

One serious problem with this approach however, is that it tends to appeal primarily to such intuitive judgments. But many people are mistakenly led by their intuitive judgments about breaking a rule to conclude that no rule was in fact broken at all. This phenomenon is known as *excuse validation* (Turri and Blouw 2015; Turri 2019). When people are given a case of someone's unintentional rule-breaking, and are asked whether they are blameworthy or criticizable for acting in violation of the rule, they almost unanimously disagree, insisting that they are blameless in such cases and so not criticizable. But when also asked if those people broke a rule, roughly half of participants *deny* that a rule has actually been broken, contrary to the set-up of the case.²⁷ Similar answers are given for cases where someone unknowingly but reason-

plicature; Kelp 2020 argues there is a knowledge norm of blaming. I do not have space here to consider the merits of these ideas.

²⁶Schechter 2018 argues against the knowledge norm largely by appealing to normative judgments about cases where one intuitively may assert without knowledge, and he gestures hopefully at a broadly Gricean account of communicative norms which abandons the idea of a norm specific to assertion; unfortunately, that approach has been considerably complicated by arguments about Gricean quality (Benton 2016b).

²⁷This result persists whether the consequences of breaking the rule are trivial or momentous: Turri 2016b, 46–47.

ably makes a false assertion. And yet, when the questions are adjusted to ask whether the person *unintentionally* broke the rule (or *unintentionally* made a false statement), the participants uniformly agree that yes, they unintentionally broke the rule (made a false statement). The explanation of these results appears to be that roughly half of participants are hesitant to claim that someone in such a case broke a rule, and so when asked if they did so, they answer “no” to avoid indirectly blaming a blameless agent; but when asked whether they unintentionally broke the rule, the explicit reference to “unintentionally” doing so frees the participant to answer correctly, because there is no risk now of indirectly blaming them. Thus presence of an excuse leads many people to mistakenly validate the agent as having broken no rule at all. As such, critics of the KNA who appeal primarily to comparable intuitive judgments about excuse and blame as a guide to what the norm of assertion must be are likely engaged in excuse validation.²⁸

Relatedly, some philosophers have objected that, intuitively, asserting in Gettier cases is perfectly permissible, and that this judgment is probative despite KNA theorists’ preferred explanation of such judgments as tracking reasonability (secondary propriety) rather than permissibility given the norm (e.g. Lackey 2007, 596–597). And of course, philosophers tend to judge most Gettier cases (or similarly structured cases) to be ones where the subject lacks knowledge. The objectors claim that these are strong counterexamples to KNA. But again, such objections do little to explain the wide range of data which KNA best explains. Moreover, experimental testing of these judgments reveal that most non-philosophers, even in Gettier cases, strongly link their judgments of knowledge with assertability, and of non-assertability with not knowing (Turri 2015; 2016b, 43–44; 2016a). This suggests, on the one hand, that ordinary speakers may regard Gettier cases as being cases of knowledge, but on the other hand, that they already implicitly accept KNA.

Finally, some philosophers complain that there is no distinctively epistemic norm of assertion, or no such stable norm which enables us to understand something about assertion. Sometimes this is because they suppose that assertion will be governed by a variety of other normative dimensions such as prudence, etiquette, morality, etc., and they speculate that these will

²⁸Cf. Kelp and Simion 2017 for a sophisticated reply to these sorts of worries.

suffice for explaining any seemingly epistemic norms on asserting (e.g. Pagin 2016, 205). Others insist that there is no well-defined category of assertion, or that it is theoretically uninteresting, or both, urging that we instead focus our theorizing on the broader notion of “sayings” (Cappelen 2011, and 2020).

These complaints do not, in my view, even amount to objections to the KNA, or to other epistemic norm views. For, first, it is no objection to a view offering the best explanation of some data merely to speculate that they might be somehow explained in other ways; one would have to articulate those explanations and argue for their superiority. And second, it is likewise no objection to KNA to insist that assertion is, to some, an uninteresting linguistic category, or one which we do not “need in order to pick out any significant component of our linguistic practice” (Cappelen 2011, 20). For those who do not find assertion a theoretically useful category, they are free to use different notions if they want to play the game of explaining the data; but if they do not even want to play that game, it is entirely unclear why they would offer their views on who is winning that game.²⁹

4 Conclusion

We’ve surveyed the arguments for the idea that knowledge is the central epistemic norm governing assertion, while nevertheless remaining neutral on the modal or other status of that norm as it figures in a broader account of the nature of assertion. KNA offers the best explanation of data from

²⁹Cappelen’s latest version of this (2020) seems to conflate locutionary acts (“sayings”) and illocutionary acts (see his “S”, p. 140), going so far as to deem the latter unworthy of disciplined theorizing. At times he suggests that his broader category of sayings is not in the business of explaining any of the relevant linguistic data, because it simply posits an act type whose tokens are subject to variable norms of prudence, etiquette, practicality, etc. (though this has been duly studied, e.g. by Turri 2017a); and he worries that “pro-assertion” views like KNA get rather complicated in their explanations of proposed counterexamples (such as some mentioned above). Yet other times, he insists that his sayings view is simpler and “can explain all the data that the pro-assertion views [like KNA] try to account for... [because they] focus on a subset of saying” (Cappelen 2020, 145). But he never aims to provide a systemic account of the norms which might explain any data supporting KNA; nor does he grasp that, if his sayings approach could explain the data covered here, it would need to invoke norms or themes specific to the illocutionary act of assertion, in which case he’d be entering into the game of explaining why the data are as they are.

several faulty conjunctions; from conversational patterns involving prompts, abstentions, challenges, hedges, parentheticals, and mutual reasoning; from hypological judgments about permission, excuse, and blame; and it also extends its insights to explanations of data from knowing how and showing how. KNA not only elegantly explains each strand; it offers a unified explanation appealing to the same knowledge-theoretic resources throughout. We've also briefly examined a few objections and complaints, finding them wanting. The case is overwhelming: knowledge is the norm of assertion.³⁰

References

- Adler, Jonathan. 2009. "Another Argument for the Knowledge Norm." *Analysis* 69: 407–411.
- Austin, J. L. 1961. *Philosophical Papers*. Oxford: Clarendon Press.
- Benton, Matthew A. 2011. "Two More for the Knowledge Account of Assertion." *Analysis* 71: 684–687.
- Benton, Matthew A. 2016a. "Expert Opinion and Second-Hand Knowledge." *Philosophy and Phenomenological Research* 92: 492–508.
- Benton, Matthew A. 2016b. "Gricean Quality." *Noûs* 50: 689–703.
- Benton, Matthew A. 2019. "Lying, Belief, and Knowledge." In Jörg Meibauer (ed.), *The Oxford Handbook of Lying*, 120–133. Oxford: Oxford University Press.
- Benton, Matthew A. 2021. "Knowledge, Hope, and Fallibilism." *Synthese* 198: 1673–1689.
- Benton, Matthew A. and van Elswyk, Peter. 2020. "Hedged Assertion." In Sanford Goldberg (ed.), *The Oxford Handbook of Assertion*, 245–263. Oxford: Oxford University Press.
- Blaauw, Martijn. 2012. "Reinforcing the Knowledge Account of Assertion." *Analysis* 72: 105–108.
- Brown, Jessica and Cappelen, Herman (eds.). 2011. *Assertion: New Philosophical Essays*. Oxford: Oxford University Press.

³⁰Many thanks to John Turri and Tim Williamson for helpful feedback.

- Buckwalter, Wesley and Turri, John. 2014. "Telling, Showing, and Knowing: A Unified Theory of Pedagogical Norms." *Analysis* 74: 16–20.
- Cappelen, Herman. 2011. "Against Assertion." In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 21–47. Oxford University Press.
- Cappelen, Herman. 2020. "Assertion: A Defective Theoretical Category." In Sanford C. Goldberg (ed.), *The Oxford Handbook of Assertion*, 139–155. Oxford: Oxford University Press.
- Carter, J. Adam, Gordon, Emma C., and Jarvis, Benjamin W. (eds.). 2018. *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- DeRose, Keith. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.
- Douven, Igor. 2006. "Assertion, Knowledge, and Rational Credibility." *The Philosophical Review* 115: 449–485.
- Green, Adam. 2017. "An Epistemic Norm for Implicature." *Journal of Philosophy* 114: 381–391.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Haziza, Eliran. forthcoming. "Assertion, Implicature, and Iterated Knowledge." *Ergo* .
- Kelp, Christoph. 2018. "Assertion: A Function First Account." *Noûs* 52: 411–442.
- Kelp, Christoph. 2020. "The Knowledge Norm of Blaming." *Analysis* 80: 256–261.
- Kelp, Christoph and Simion, Mona. 2017. "Criticism and Blame in Action and Assertion." *Journal of Philosophy* 114: 76–93.
- Kvanvig, Jonathan L. 2009. "Assertion, Knowledge, and Lotteries." In Patrick Greenough and Duncan Pritchard (eds.), *Williamson on Knowledge*, 140–160. Oxford: Oxford University Press.
- Lackey, Jennifer. 2007. "Norms of Assertion." *Noûs* 41: 594–626.
- Lackey, Jennifer. 2011. "Assertion and Isolated Second-Hand Knowledge." In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 251–275. Oxford: Oxford University Press.
- Lackey, Jennifer. 2016. "Assertion and Expertise." *Philosophy and Phenomenological Research* 92: 509–517.
- MacFarlane, John. 2011. "What is Assertion?" In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 79–96. Oxford: Oxford University Press.

- MacIver, A.M. 1938. "Some Questions about 'Know' and 'Think'." *Analysis* 5: 43–50.
- Maitra, Ishani. 2011. "Assertion, Norms, and Games." In Brown and Cappelen (2011).
- McKinnon, Rachel. 2013. "The Supportive Reasons Norm of Assertion." *American Philosophical Quarterly* 50: 121–135.
- Moore, G.E. 1962. *Commonplace Book: 1919–1953*. London: George Allen & Unwin.
- Pagin, Peter. 2016. "Problems with Norms of Assertion." *Philosophy and Phenomenological Research* 93: 178–207.
- Papineau, David. 2021. "The Disvalue of Knowledge." *Synthese* 198: 5311–5332.
- Schechter, Joshua. 2018. "No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion." In Carter, Gordon, and Jarvis 2018, 132–159.
- Simion, Mona and Kelp, Christoph. 2020. "Assertion: The Constitutive Norms View." In Sanford C. Goldberg (ed.), *The Oxford Handbook of Assertion*, 59–73. Oxford: Oxford University Press.
- Slote, Michael. 1979. "Assertion and Belief." In Jonathan Dancy (ed.), *Papers on Language and Logic*, 177–191. Keele: Keele University Library. Reprinted in Slote (2010): 94–102.
- Slote, Michael. 2010. *Selected Essays*. New York: Oxford University Press.
- Sosa, Ernest. 2010. "Value Matters in Epistemology." *Journal of Philosophy* 107: 167–190.
- Sosa, Ernest. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Turri, John. 2010. "Prompting Challenges." *Analysis* 70: 456–462.
- Turri, John. 2011. "The Express Knowledge Account of Assertion." *Australasian Journal of Philosophy* 89: 37–45.
- Turri, John. 2014. "Knowledge and Suberogatory Assertion." *Philosophical Studies* 167: 557–567.
- Turri, John. 2015. "Knowledge and the Norm of Assertion: A Simple Test." *Synthese* 192: 385–392.
- Turri, John. 2016a. "Knowledge and Assertion in 'Gettier' Cases." *Philosophical Psychology* 29: 759–775.
- Turri, John. 2016b. *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*. Cambridge: Open Book Publishers.

- Turri, John. 2016c. "The Point of Assertion is to Transmit Knowledge." *Analysis* 76: 130–136.
- Turri, John. 2016d. "Vision, knowledge, and assertion." *Consciousness and Cognition* 41: 41–49.
- Turri, John. 2017a. "The Distinctive 'Should' of Assertability." *Philosophical Psychology* 30: 481–489.
- Turri, John. 2017b. "Experimental Work on the Norms of Assertion." *Philosophy Compass* 12: 1–9.
- Turri, John. 2018. "Sustaining Rules: A Model and Application." In Carter, Gordon, and Jarvis 2018, 259–277.
- Turri, John. 2019. "Excuse Validation: A Cross-Cultural Study." *Cognitive Science* 43: e27148.
- Turri, John and Blouw, Peter. 2015. "Excuse Validation: A Study in Rule-Breaking." *Philosophical Studies* 172: 615–634.
- Unger, Peter. 1975. *Ignorance: A Defense of Skepticism*. Oxford: Clarendon Press.
- van Elswyk, Peter. 2021. "Representing Knowledge." *The Philosophical Review* 130: 97–143.
- Weiner, Matthew. 2005. "Must We Know What We Say?" *The Philosophical Review* 114: 227–251.
- Williamson, Timothy. 1996. "Knowing and Asserting." *The Philosophical Review* 105: 489–523.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. forthcoming. "Justification, Excuses, and Sceptical Scenarios." In Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon*. Oxford: Oxford University Press.