Should I say that? An experimental investigation of the norm of assertion

Abstract: Assertions are our standard communicative devices for sharing and acquiring information. Recent studies seemingly provide converging evidence that assertions are subject to a *factive* norm: you are entitled to make an assertion only if it is true. However, these studies assume that we can treat participants' judgements about what an agent 'should say' as evidence of their intuitions about assertability. This paper argues that this assumption is incorrect, so the conclusions drawn in the aforementioned studies are unwarranted. We provide evidence that most people do not interpret statements about what one 'should say' as statements about assertability, but rather as statements about what is in the agent's interest to do. Measures for prompting the intended reading of statements about what an agent 'should say' are identified, and their efficacy is tested. We found that when these measures are implemented, people's judgements consistently and overwhelmingly align with non-factive accounts of assertion.

Keywords: communication, truth, linguistic normativity, intuitions, rules.

1. Introduction

1.1 The norm of assertion

Communication is a fundamental feature of human life, and our collective well-being depends significantly on our linguistic ability to share reliable information about the world that we inhabit. Ordinarily, we share information by making *assertions* – that is, by claiming that something is the case. Since assertions are so important for sharing information, it is not surprising that they are constrained by some cognitive and epistemic expectations. We count on each other not to make insincere claims, and we typically criticize other speakers if we find out that they lied, or that their statements were not based on adequate evidence. But what exactly do we expect from our fellow communicators? Or, to put it more precisely: under which conditions do we take other speakers to be epistemically entitled to make an assertion?

Over the past two decades, this fundamental question about human communication has taken centre stage in philosophy of language and epistemology. A large body of literature deals with a hypothesis first put forward by Williamson (1996), according to which assertion is regulated by a single norm, which enjoins speakers to assert only those propositions that meet a certain

epistemic threshold. Williamson proposes to formalize the envisaged rule as follows: “one must: assert that p only if p has C ”, where C is an epistemic property of the asserted proposition.

To illustrate, suppose that C is the property of *being true*. It would follow that one is entitled to assert a proposition p only if p is true – so that if p is false and you assert it, you are violating the norm. But Williamson’s hypothesis leaves open the possibility that the threshold for proper assertion could be something other than truth: for instance, C could be the property of *being known by the speaker*, or simply the property of *being believed*. If we accept Williamson’s hypothesis¹, understanding the cognitive and epistemic expectations governing human communication crucially involves determining which epistemic standard ‘ C ’ makes a proposition assertable.

1.2 Competing accounts of the norm of assertion

What is the norm of assertion? That is, what is the property ‘ C ’ that makes a proposition assertable? Many different proposals have emerged in the scholarly debate, the most influential of which are the following:²

- KNOWLEDGE RULE: “Assert p only if you know that p ”
(DeRose, 2002; Engel, 2008; Hawthorne, 2004; Reynolds, 2002; Williamson, 1996)
- TRUTH RULE: “Assert p only if p is true”
(Alston, 2000; MacFarlane, 2014; Weiner, 2005; Whiting, 2012)
- JUSTIFICATION RULE: “Assert p only if you rationally believe that p ” (or “only if it is rational for you to believe that p ”)
(Doven, 2006; Gerken, 2012, 2017; Kvanvig, 2009; Lackey, 2007; McKinnon, 2012, 2013)
- BELIEF RULE: “Assert p only if you believe that p ”
(Bach, 2008; Hindriks, 2007)

¹ For some critical takes on the hypothesis, see Cappelen (2011), Carter & Gordon (2011), Pagin (2016), Carter (2017), Marsili (2019). While we share a critical stance on these matters, we also think that (despite its weaknesses) Williamson’s hypothesis can be a useful *idealization* to guide the investigation of these issues.

² This overview is inevitably not exhaustive. Given the vast amount of literature on this subject, we left out some alternative formulations of each position (such as the knowledge-provision rule, as defended by García-Carpintero, 2004, and Pelling, 2013, or the context-sensitive proposal defended by Goldberg, 2015) and brushed over some nuances, grouping together views that differ in some important respects (such as the different accounts of epistemic justification and rationality that are deployed to characterize the justification rule). For a more detailed review, see Pagin (2014).

One of the key points of disagreement between these views is whether speakers are only entitled to make assertions that are in fact *true*. To put it more succinctly, the question is whether assertability requires truth (whether *C* entails that *p* is true). An example (from Turri, 2013, p. 282) may help to illustrate where each view stands on this issue:

ROLEX UNLUCKY

Maria is a watch collector. She owns so many watches that she cannot keep track of them all by memory alone. So she maintains a detailed inventory of them. She keeps the inventory up to date. Maria knows that the inventory is not perfect, but it is extremely accurate. Today Maria is having guests over for dinner.

Soon after dinner is served, one of her guests asks, “Maria, do you have a 1990 Rolex Submariner in your watch collection?”

Maria consults her inventory. It says that she does have a 1990 Rolex Submariner in her collection. But this is one of those rare cases where the inventory is wrong: she does not have one.

Suppose that in this scenario Maria replies: “Yes, I have a 1990 Rolex Submariner in my collection.” Here Maria says what she believes to be true, but she ends up saying something false. We may call her assertion ‘unlucky’, to indicate that its falsity is due to bad luck, rather than bad faith. Maria’s unlucky assertion allows us to illustrate where each view stands with respect to the assertability of false propositions. The verdict of the knowledge rule³ and the truth rule is that her assertion is inappropriate: an assertion is permissible only if it is true, so Maria is not entitled to make her assertion. We may call these views *factive*, to indicate that they require speakers to state only the facts. The belief rule and the justification rule, by contrast, deem Maria’s assertion permissible: they are *non-factive*, in that they do not require speakers to state only the facts. To summarize:

FACTIVE RULES: Knowledge rule, Truth rule

NON-FACTIVE RULES: Justification rule, Belief rule

Both factive and non-factive accounts are typically motivated by observations about our linguistic intuitions and behaviour. Advocates of non-factive rules tend to stress that unlucky assertions are intuitively permissible: since we would not deem Maria’s response improper or criticizable, a good account of assertion should predict that she is not violating any norm. To be

³ Knowledge is understood to entail *belief*, *justification* and *truth*, so that in order to follow the knowledge rule, you also have to follow all the other rules, including the truth rule.

sure, non-factivist philosophers acknowledge that unlucky assertions are sub-optimal. They contend, nonetheless, that unlucky assertions do not violate any communicative norm: in slogan form, their view is that ‘assertability does not require truth’.

Proponents of factive accounts tend to challenge this intuition: they argue that Maria is not entitled to assert that she has a Rolex 1990 Submariner, given that she does not have one. More generally, factive accounts aim to make sense of the rather straightforward intuition that, *ceteris paribus*, a false assertion is improper and criticizable. After all, we criticize people for making false assertions (and not for making true ones), and we feel at fault when we discover that we have said something false. Falsity seems to constitute a distinctive kind of wrongness for assertions, and factive views are well equipped to identify the normative source of this wrongness: false assertions violate factive norms of assertion.

Settling the disagreement between factive and non-factive accounts of the norm of assertion is not easy. The philosophical literature is ripe with theoretical (non-empirical) arguments supporting various species of factive and non-factive views (for an overview, see Pagin, 2014). This paper will not deal with such arguments: it will rather be concerned with whether there is *experimental* support for one view and against the other. In this respect, we follow some authors (Douven, 2006, p. 450; Pagin, 2016, p. 22; Turri, 2013) who think that the norm of assertion hypothesis is essentially an empirical hypothesis, and that disagreement about the norm of assertion can be settled (or at least significantly advanced) by verifying which account is best supported by the available evidence (see Knobe & Nichols, 2013, for how empirical approaches can advance philosophical discussions).

1.3 Experimental research on the norm of assertion

To determine which rules govern natural languages, researchers in psycholinguistics typically appeal to the behaviours and intuitions of competent speakers. Competent speakers are often said to possess *procedural*, as opposed to *explicit* (or *declarative*), understanding of the relevant rules: they are able to follow them, even if in many cases they lack an explicit understanding of their nature (Chomsky, 1965, pp. 4–8; Mikhail, 2011, pp. 41–42; Searle, 1969, pp. 41–42). With appropriately designed experimental tasks, researchers can prompt competent speakers to translate their procedural competence into action. This allows them to identify patterns of behaviour, and determine which theory best fits those patterns. This approach has been adopted across a variety of disciplines studying human languages, to reconstruct the syntactic (e.g.

Chomsky, 1965), and semantic/pragmatic (e.g. Musolino, 2009; Noveck, 2001; Noveck & Sperber, 2004) norms that govern our conversational exchanges⁴.

The same approach has been used to empirically investigate the norm of assertion. Turri (2013), for instance, invited participants to read the ROLEX UNLUCKY vignette reported above. When presented with the test question “Should Maria tell her guest that she has a 1990 Rolex Submariner in her collection?” [Yes/No], subjects overwhelmingly selected the latter option (No), even when different factors were manipulated (the control questions, the stakes, and the response options), providing apparently robust evidential support for factive accounts. Other studies found similar results, accumulating an impressive body of evidence supporting factive accounts of assertion in general (Turri, 2013, 2017b, 2018, 2020) and the knowledge rule in particular (Turri, 2014a, 2014b, 2015, 2016; Turri, Friedman & Keefner, 2017; Turri & Park, 2018; Turri & Buckwalter, 2017; Turri, 2018).

It seemed that the debate had been settled in favour of factive norms, until a few years ago new studies came out that pointed in the opposite direction, suggesting that the norm governing assertion was instead non-factive, and modellable along the lines of the justification rule (Kneer, 2018; Reuter & Brössel, 2019). Although some of the claims made in these studies have since been challenged (Turri, 2018, 2020), the case for factive accounts is now less straightforward: the emergence of evidence pointing in the opposite direction calls for an explanation.

The primary goal of this paper is to provide such an explanation. We identify a potential flaw in the studies supporting factive accounts: a problematic ambiguity in how the test question is phrased. If we are right, the main test question of these studies was not interpreted in a way that is relevant to test the norm of assertion hypothesis. Our Experiment 1 finds support for this criticism. Some measures are then identified to minimise the ambiguity of the test question, and to check whether participants really interpreted it as intended. Experiment 2 and 3 show that when these measures are implemented, competent speakers systematically judge unlucky assertions to be permissible. Based on these findings, we conclude that non-factive accounts are in a better position to accommodate the available empirical data.

1.4 A problem with studies that support factive norms

Evidence for a factive norm is apparently very robust, as it comes from a variety of studies (Turri, 2013, 2014a, 2014b, 2015, 2016; Turri, Friedman & Keefner, 2017; Turri & Park, 2018; Turri & Buckwalter, 2017; Turri, 2017a, 2017b, 2018, 2020) whose results converge irrespectively of the

⁴ See Gibson & Fedorenko (2013) on the need for quantitative methods in linguistics.

vignettes adopted, the demographics of participants, the design of the experiment, and so forth. Crucially, however, in all their differences these studies share a central methodological aspect: they explore laypeople’s intuitions by asking a sample of subjects to judge what a particular agent *should do* in a given scenario. Let us focus on studies of unlucky assertions (Turri, 2013, 2017b, 2018; Turri & Park, 2018; Turri 2020), which are directly relevant to the question animating this paper (whether the norm of assertion is factive or non-factive). In the *experimentum crucis* of these studies, participants are presented with a vignette in which the protagonist is about to make an unlucky assertion, and are prompted to judge whether the protagonist should make that assertion.

For a concrete example, let us consider again Exp.1 from Turri (2013). Here participants had to read the ROLEX UNLUCKY vignette and then answer the question: “Should Maria tell her guest that she has a 1990 Rolex Submariner in her collection?” [Yes/No]. The underlying idea is that participants will answer “Yes” if they deem unlucky assertions to be assertable, and “No” if they do not. The studies conducted by Turri and colleagues consistently employed an experimental design along these lines: intuitions about assertability are investigated almost exclusively by asking participants to judge whether the protagonist of a vignette *should* make a given statement. Throughout these studies, knowledge proves to be a better predictor of assertability judgements than any of its rivals (certainty, justification, belief – except for Turri, 2017b, where truth does better than knowledge).

These studies thus draw conclusions about assertability from laypeople’s judgements about whether a given agent ‘should’ assert something. They all rely on a crucial assumption: that if a participant judges that an agent should assert *p*, then she judges *p* to be assertable, and that if a participant judges that an agent should not assert *p*, then she judges *p* not to be assertable. Call this assumption the *assertability* assumption. Turri (2013) offers a plausible justification for the assertability assumption, and recognizes its limits:

The literature on assertion’s norm suffers from some terminological inconsistency. What is the right way to express assertion’s constitutive norm? Some say, ‘You should make an assertion only if ...’; others say ‘You ought to make an assertion only if ...’; others say, ‘You may make an assertion only if ...’; and others say, ‘You must: make an assertion only if ...’. [...] I’m not going to resolve this inconsistency here. Instead, I will simply opt for the ‘should’ formulation. And when probing laypeople I will stick to asking about what a speaker ‘should’ say or whether anything ‘incorrect’ has been done. It is legitimate to question whether different terminology would lead to different results. But choices must be made in order to get the project off the ground and I have chosen to start here. I welcome and encourage further work that makes different choices. (2013: 281)

However, subsequent studies by Turri continue to employ ‘should-questions’ as the standard test for assertability. This means that the significance of these studies is conditional on the validity of

the assertability assumption: if the assertability assumption is right, then the evidence collected by these studies supports factive accounts (and the knowledge norm in particular). But if the assertability assumption is wrong, then the empirical argument for factive accounts (and the knowledge norm) is inconclusive.⁵

This paper starts from a simple observation that puts some pressure on the assertability assumption. We take it that there are two natural readings of the verb ‘should’ that was employed in the test questions. On the one hand, we may say that an agent ‘should do something’ *in order to follow a norm* – this is what we call a *deontological* reading. On the other hand, we may say that an agent ‘should’ do something *in order to meet their aims, or some other standard for success* – this is what we call a *teleological* reading. Only the first reading is compatible with the assertability assumption: only if should-questions are interpreted deontologically (rather than teleologically) can they be evidence of laypeople’s intuitions about the *norm* of assertion, as opposed to evidence of their intuitions about what makes an assertion *successful*.

1A

1B

Figures 1A and 1B: Connect-four positions illustrating two different senses of ‘should’.

⁵ Our point here is that these studies *need* to make this assumption, not that they *explicitly acknowledge* that such an assumption is needed. Without the assertability assumption (or one along its lines) these studies would be unable to achieve their stated aim, namely to collect evidence that can help to settle the disagreement about the norm of assertion.

To illustrate this point, let us consider how the term ‘should’ can be used in two different and incompatible ways in a different context: that of describing what a player of Connect Four⁶ *should* do. Figure 1A (left) shows Red’s turn, in which Red has to move. All columns but one are fully occupied, so Red can only drop his disc in the seventh (rightmost) column. This is not a convenient move for Red, since it will allow Yellow to put his disc in the top-right corner, thus winning the game by connecting four Yellow discs diagonally. Although it would be more convenient for Red to skip his turn, in this situation Red has to drop his disc in the last column, because the rules of Connect Four require each player to drop a disc during their turn. In this situation we would say that Red *should* drop his disc in the seventh column, as no other legal move is available. Here ‘should’ indicates what Red is supposed to do in order to *follow the rules* of the game: it is a ‘should’ with a *deontological* value.⁷

But now consider the *previous* turn (Figure 1B, right; Yellow’s turn). In this position, Yellow has two options: he could drop his disc in the second or the seventh column. It is equally natural and correct to say that Yellow *should* drop his disc in the second column, because this is the only move that will lead him to win the game – it is the move that will lead to Figure 1A, in which Red will be forced to make his losing move. Here ‘should’ does not indicate what Yellow should do in order to follow the rules of Connect Four (another legal move is available, namely dropping the disc in the seventh column), but rather what Yellow should do in order to achieve the aim of winning the game: it is a *teleological* should. This illustrates the two senses in which the term ‘should’ is used in ordinary language: this verb has a ‘deontological’ value when it indicates which course of action is required *to follow a rule*, and a ‘teleological’ value when it indicates which course of action is required or optimal *to meet an aim*.

Two points are worth stressing in relation to this. First, teleological ‘shoulds’ are not typically reducible to deontological ‘shoulds’. When we say that Yellow should drop his disc in the second column (in 1B), ‘should’ has *only* a teleological value: it clearly does not mean “should, in order to follow the rules of the game”. The teleological ‘should’ here identifies a move that, within a range of legal moves, is optimal for meeting a success condition (winning the game). Hence,

⁶ For the uninitiated, the aim of the game here is to form a horizontal, vertical, or diagonal line of four discs of your own colour. During each turn a player must drop a disc in one of the seven columns, where the disc will fall, occupying the lowest space available.

⁷ Figure 1A is a ‘zugzwang’ position: these are positions (in games that compel each player to move during their turn) in which every available move will lead the moving player to a significant loss (Golombek, 1977). In zugzwang positions the player “should” make a move in the deontological sense, but clearly not in the teleological sense, for each move is against their interest.

there are contexts in which ‘should’ is interpreted correctly *only* in the teleological sense.⁸ Second, both uses of the verb are common and natural. There is nothing ‘odd’ or ‘artificial’ in saying that (in 1A) Red *should* drop his disc in the seventh column (in the deontological sense), and similarly there is nothing off about saying that (in 1B) Yellow should drop his disc in the second column (in the teleological sense).⁹

If ‘should’ is ambiguous between a teleological and a deontological interpretation, then the evidence collected by Turri and colleagues is inconclusive, since their test questions could have been interpreted in two different and incompatible ways, only one of which supports their conclusions (on unwanted interpretations of research questions, see Schwarz, 2014; Royzman & Hagan, 2017; Wiegmann, Waldmann & Samland, 2016).¹⁰ To see this, consider once again the previous example from Turri (2013). Here participants could have interpreted the test question (“Should Maria tell her guest that she has a 1990 Rolex Submariner in her collection?”) deontologically (as a question about whether Maria violated a rule) or teleologically (as a question about whether Maria failed to meet some aim or standard for success). Only a prevalence of deontological interpretations would be compatible with the assertability assumption, thus supporting Turri’s conclusions about laypeople’s preference for factive views. A prevalence of teleological interpretations, by contrast, would undermine such conclusions: it would mean that the participants expressed judgements that are irrelevant to the existing debate on the norm of assertion. If the participants indicated that in saying something false Maria *fails to fulfil her aims*, rather *than failing to comply with a linguistic rule*, then we are not licensed to infer that they took her

⁸ For more general discussions of the distinction between rules (deontological normativity) and aims (teleological normativity) in relation to games in general and assertion in particular, see Kemp (2007) and Marsili (2018).

⁹ Admittedly, the teleological ‘should’ is sometimes more natural to employ than its deontological counterpart – but this only puts further pressure on the assertability assumption. We think that there is an explanation for this slight difference. When a speaker can choose between a less informative and a more informative expression, employing the less informative one is typically perceived as odd or uncooperative (see e.g. Grice, 1989, and Levinson, 1983, on the Quality Maxim and scalar implicatures). This may explain why the teleological reading is more natural: the deontological ‘should’ allows for alternatives that are not ambiguous between these two interpretations (like ‘must’, ‘have to’, or ‘is obliged to’). By contrast, the teleological ‘should’ lacks unambiguous alternatives of this kind.

¹⁰ Turri has attempted to address a similar (but unrelated) objection. He designed an experiment to test the hypothesis that “the ‘should’ of assertability is essentially tied to whether the assertion is true or known by the speaker” against the hypothesis that it is tied instead to other forms of normativity, “such as morality, practical rationality, etiquette, and legality” (2017b, p. 486). The study found that “evaluations of truth value were the strongest predictor” (ibid.). Crucially, this study does not address the distinction between teleological and deontological readings of the verb ‘should’, which is what we take to undermine the assertability assumption. To show that truth is a better predictor than practical rationality, etiquette and legality is not to show that ‘should questions’ are interpreted as questions about assertability, which is what one would need to prove in order to support the assertability assumption.

false statement to violate the norm of assertion. The aim of our first experiment is to verify how participants interpreted the should-questions in Turri's (2013) "test of truth", in order to assess whether this test does indeed track intuitions about the norm of assertion, or intuitions about some other standard of evaluation instead.

2. Experiment 1 – Follow-up questions

In order to test how participants interpreted the test question, we replicated the crucial condition of the "test of truth" (Turri, 2013), ROLEX UNLUCKY, and added a follow-up question to check how participants who gave factive answers interpreted the should-question. The follow-up question allowed us to test whether these participants interpreted the verb 'should' teleologically or deontologically. If the majority of participants adopt a deontological interpretation, then there is little reason to doubt that the assertability assumption is correct, and Turri's results actually support the factive view. But if most participants adopt a teleological interpretation, then the assertability assumption is incorrect, and this method of inquiry fails to test intuitions about the norm of assertion; it rather tests intuitions about what is optimal for the protagonist to do, given her goals.

Participants

In all the experiments reported in this paper, participants were recruited on Prolific Academic (Palan & Schitter, 2017) to complete an online survey implemented on the Unipark online platform. We recruited only adult native English speakers, with an approval rate on the Prolific Platform of at least 90%. All our experiments were preregistered on OSF (see Appendix).

We excluded participants if they met one or more of the following preregistered conditions: failing the attention check, failing the control questions, or taking less than 40 seconds to complete the survey. As a result, of the 279 participants who started the survey (71% female, age $M=33$ years), 202 were included in the study. Participants received £0.20 for an estimated two minutes of their time (£6/h).

Design, materials, and procedure

Participants first read general instructions to familiarize them with the task and the response formats. They were then randomly assigned to one of two conditions (test question: EXCLUSIVE vs. INCLUSIVE), between-subjects design. They were all presented with the original ROLEX

UNLUCKY ASSERTION scenario (see §2), followed by the original control questions (Q1 and Q2) and the original test question (Q3):

Q1: Is there a 1990 Rolex Submariner in Maria's collection? (yes/no)

Q2: If Maria tells her guest that she has a 1990 Rolex Submariner in her collection, she will be saying something... (true/false)

Q3: Should Maria tell her guest that she has a 1990 Rolex Submariner in her collection? (yes/no)

Since this experimental setup was virtually identical to the one employed in Turri (2013), we predicted that the majority of participants would answer 'no' to Q3. To understand which interpretation of the test question (teleological or deontological) underlay these answers, all and only participants who answered 'no' to Q3 (N=180) were redirected to the follow-up page, and randomly assigned to one of the following two conditions (which differ in how the follow-up question is formulated). Response options were always ordered randomly.

In the INCLUSIVE condition (see below), participants were allowed to select any number of explanations (i.e. zero up to three) for their initial judgement. They could select a teleological explanation (Maria fails to meet an aim), a deontological explanation (Maria fails to follow a rule), or a non-epistemic explanation (Maria's assertion harms the guest). The EXCLUSIVE condition (also below) involved a sentence-completion task, and participants had to select one of two explanations: a teleological one, or a deontological one.¹¹ Since our hypothesis was that in this experimental setup it is natural to interpret 'should' teleologically, we predicted that the majority of participants (significantly more than 50%) would choose the teleological option in both conditions.

Your judgment indicates that you think that Maria should not tell her guest that she has a 1990 Rolex Submariner in her collection. Please explain your judgment by [(A): picking the options that explain what you meant by "should not"/ (B): choosing the option that best expresses your intuition]

(A): INCLUSIVE condition

(You can pick none and up to three options.)

- She "should not" do that, otherwise she will fail in her intention to tell the truth
- She "should not" do that, otherwise she will violate the norms of conversation

¹¹ Each condition had an experimental advantage. The INCLUSIVE condition allowed participants to reject the deontological/teleological distinction altogether, in a number of ways: by picking both options, no option, or the third option. The EXCLUSIVE condition lacked this freedom of choice, and forced subjects to express a preference for one interpretation over the other: a preference that we aimed to ascertain, but that the INCLUSIVE condition allowed them not to express.

- She “should not” do that, otherwise her guest will face bad consequences

(B): EXCLUSIVE condition

(You can and have to pick only one option.)

- ...because otherwise she would fail in her intention to tell the truth – even if she would not violate a rule
- ...because otherwise she would fail in her intention to tell the truth, and she would also violate a rule

After completing this task, participants were directed to a new screen and asked to respond to some demographic questions, and a simple transitivity task to test for attention.

Figure 2: On the left, percentages of participants choosing the teleological vs. deontological option (forced-choice) in the EXCLUSIVE condition. On the right, percentages of participants picking vs. not picking each of the three options (teleological, deontological, non-epistemic) in the INCLUSIVE condition (where picking any number of options, including no option, was allowed). Error bars represent 95% confidence intervals.

Results

All our experimental findings were consistent with our preregistered predictions (see Appendix). The majority of participants (180 out of 202: 89%) denied that Maria should say that she has the Rolex (factive answer) (binomial test against 50%, $p < .001$; two-tailed), replicating the results of Turri (2013). Of the 180 participants who gave factive responses, the majority indicated that they understood the should-question teleologically (see Fig. 2) in both follow-up conditions. In the INCLUSIVE condition (where up to three options could be selected), 91% (82/90) picked the teleological option (binomial test against 50%, $p < .001$; two-tailed), 10% (9/90) the deontological

one, and 2% (2/90) indicated a consequentialist explanation. The teleological interpretation was selected more often than both the factive option, $\chi^2_{1, N=180} = 118.44, p < .001$, and the consequentialist option, $\chi^2_{1, N=180} = 142.86, p < .001$. In the EXCLUSIVE condition, the majority (64%; 58/90) chose the teleological options (binomial test against 50%, $p = .008$, two-tailed).

Discussion

These results challenge a fundamental premise underlying Turri's (2013) study: the assertability assumption, according to which should-judgements track judgements of assertability. Participants overwhelmingly interpreted the test question teleologically (as a question about what is in Maria's interest to do), rather than deontologically (as a question about what Maria is supposed to do), which is incompatible with this assumption. This in turn means that the data collected in Turri (2013) neither supports nor undermines factive accounts of assertion: it is simply irrelevant to the existing philosophical debate on the norm of assertion, because data about what laypeople think is *conducive to the speaker's goal* of telling the truth tells us nothing about what laypeople think a speaker is *allowed to do*.

This experiment shows, against the assertability assumption, that questions about what an agent 'should say' do not necessarily track assertability judgements, because of the ambiguity between teleological and deontological readings of the verb 'should'. This is a significant discovery, as the whole body of studies conducted by Turri and colleagues (Turri, 2013, 2014b, 2014c, 2015, 2016, 2017b, 2017a, 2018, 2020; Turri & Buckwalter, 2017; Turri et al., 2017; Turri & Park, 2018) almost invariably employs this questioning format. While we are not claiming that the participants of all these studies *must* have interpreted the task in a way that is incompatible with the assertability assumption (i.e. teleologically rather than deontologically), our results highlight that the opposite cannot be taken for granted, calling into question the conclusions drawn in these studies. That participants' should-judgements track judgements of assertability in these experiments is, at most, a conjecture. For unless one controls for how should-judgements are interpreted, there are no strong reasons to assume that participants will adopt a deontological interpretation, rather than a teleological one.

Another interesting pattern that emerges from these results is the following: the findings indicate that participants not only consistently report a teleological reading of the test question, but also reject normative ones. This is especially clear in the EXCLUSIVE condition, where most participants indicated that Maria would *not* violate a rule in saying something false. These responses are hard to square with the predictions of factive accounts. If we are to trust participants' self-reports, these results not only undermine previous evidence found in support of

factive accounts and the knowledge norm of assertion, but also represent new (circumstantial) evidence that laypeople’s intuitions rather align with non-factive accounts (in line with recent findings by Kneer, 2018, and Reuter & Brössel, 2018). And this is indeed what our next two experiments suggest: whenever participants indicate that they interpreted the test question in a deontological way, they consistently and overwhelmingly judge unlucky assertions to be permissible.

3. Experiment 2 – *Past-Should Design*

An effective “test of truth” for the norm of assertion must avoid the problematic ambiguity between deontological and teleological readings: it must track judgements of assertability, rather than other considerations about whether an assertion should be made (e.g. whether the assertion would be in the speaker’s interest). Experiment 2 aims to refine the original design in order to exclude teleological interpretations and effectively test for assertability.

We devised a simple method for attempting to make teleological interpretations of should-judgements less likely: manipulating the temporal order of the events. In the original ROLEX UNLUCKY scenario, Maria has not yet made her assertion, so it is perfectly natural to ask which course of action will lead Maria to succeed in her intention to answer the question correctly: here the *teleological* reading of the verb ‘should’ in the test question is the most relevant and salient reading. But if one modifies the scenario so that Maria has already made her false assertion, the teleological reading of the test question should be less salient: after all, it is too late to *recommend* which course of action is in Maria’s interest (teleological reading), but it is not too late to point out that her behaviour was *improper* (deontological reading). Considerations about which course of action is *preferable* are of little relevance after Maria has made her statement (it is too late for Maria to change her actions), whereas considerations about which course of action is *permitted* are still relevant (Maria may decide to apologize, and we – or a third party – may decide to criticize her, condemn her or excuse her). If these observations are right, then the deontological reading should be the most salient one when Maria has already made her assertion. To test this hypothesis, in Experiment 2 we presented participants with a modified version of the vignette, in which Maria has already spoken falsely, and investigated once again how participants interpreted the main question, using follow-ups.

Participants

We recruited 487 participants with the same method and exclusion criteria as the previous experiment, preventing previous participants from taking the survey. Data was collected until the preregistered criteria were met, resulting in 300 participants being included in the analysis (69% female, age $M=26.9$ years). Participants received £0.20 for an estimated two minutes of their time (£6/h).

Design, materials, and procedure

After reading the usual instructions, participants were randomly assigned to one of three between-subjects conditions (PRESENT-POSITIVE, PAST-POSITIVE, PAST-NEGATIVE).¹² In each condition, participants were presented with a slightly modified version of the ROLEX UNLUCKY vignette, in which it is made clearer that Maria's inventory is reliable.¹³ The PRESENT-POSITIVE condition is otherwise identical to the original vignette: Maria has not yet spoken when subjects are interrogated about what she 'should say'. In the PAST-POSITIVE and PAST-NEGATIVE conditions, by contrast, Maria has already made her false statement when participants are asked to make their judgement (cf. Appendix). The participants' main task was to agree or disagree with a statement, whose tense was different in each condition (*present tense* in PRESENT-POSITIVE, *positive past tense* in PAST-POSITIVE, *negative past tense* in PAST-NEGATIVE; underlined below, but not in the survey).

(Q1) PRESENT-POSITIVE: Maria should tell her guest that she has a Rolex 1990 Submariner.
[Agree/Disagree]

(Q1) PAST-POSITIVE: Maria should have told her guest that she has a Rolex 1990 Submariner.
[Agree/Disagree]

(Q1) PAST-NEGATIVE: Maria should not have told her guest that she has a Rolex 1990 Submariner.
[Agree/Disagree]

¹² Although we are reporting three conditions, the study had an additional one, meant to verify that our transition from "yes/no" to "agree/disagree" as response options did not affect people's response patterns. Such a transition was introduced to facilitate participants' understanding of the response options in the negatively worded condition (PAST-NEGATIVE): "Yes/no" response options would have been ambiguous in this condition, since in ordinary language both can be used to indicate agreement with the relevant claim ("yes, she should not have said that p " and "no, she should not have said that p "). Since, as expected, the change from "yes/no" to "agree/disagree" did not affect participants' responses, we follow the recommendation of an anonymous reviewer to move our analysis of this extra condition into the Appendix.

¹³ To ensure this, we removed the sentence "Maria knows that the inventory is not perfect, but it is extremely accurate" (see Appendix). We suspected that it made it unclear whether Maria was epistemically justified to trust her inventory, and to which extent. Whether it is problematic to leave the sentence in the vignette has since been discussed in the literature (Reuter & Brössel 2018 claim that it is, Turri 2020 that it is not).

In all conditions, this task was followed by two control questions:

Q2: It is reasonable for Maria to believe that she has a Rolex 1990 Submariner. [Agree/Disagree]

Q3: Maria has a Rolex 1990 Submariner. [Agree/Disagree]

The subjects in the PAST-POSITIVE and PAST-NEGATIVE conditions were presented with follow-up questions, to verify whether the test questions in the past tense were interpreted as intended. Participants were sorted depending on their response to the test question Q1 (see flowchart in Figure 3), and then directed to a new screen with follow-up questions.

Participants whose response aligned with a factive view ('disagree' with PAST-POSITIVE, 'agree' with PAST-NEGATIVE) were presented with a revised version of the follow-up completion task from experiment 1 (EXCLUSIVE condition, tense altered to match the fact that Maria has already spoken).

Q4: Maria should not have told her guest that she has a 1990 Rolex Submariner in her collection...

1. ... because otherwise she would fail in her intention to tell the truth – even if she would not violate the rules governing the conversation
2. ... because otherwise she would fail in her intention to tell the truth, and she would also violate the rules governing the conversation

This follow-up task would not be a helpful check for participants whose answers aligned with non-factive views ('agree' in PAST-POSITIVE, 'disagree' in PAST-NEGATIVE): here we could not expect a rational participant to pick the teleological explanation, as there is no reason to assume that Maria *intends* to make a false statement. The worry about these participants is rather that they could merely try to indicate that Maria's false assertion is *excusable*, rather than permissible¹⁴ (so that their intuitions align with factive accounts after all). To control for this potential ambiguity (following a method devised and tested by Turri & Blouw, 2015, pp. 625-7), we allowed participants to either clarify that Maria inadvertently violated a rule, or insist that her behaviour was permissible (the formulation varied depending on whether the participant had been redirected from PAST-POSITIVE or from PAST-NEGATIVE):

Q4: [Maria should have told her guest /it was permissible for Maria to tell her guest] that she has a 1990 Rolex Submariner in her collection...

1. ... because she violated the rules of the conversation, but did it inadvertently

¹⁴ For a detailed discussion of the hypothesis that the availability of excuses could pollute our intuitions about whether a factive rule has been violated, see Boulton (2017), Gerken (2011, sec. 3-4), Madison (2017), Schechter (2017), Turri & Blouw (2015), Williamson (2000, sec. 11.4; forth), and section 5.2 of this paper.

2. ... because she did not violate the rules of the conversation, given that she had the best reasons to believe that what she said was true.

Figure 3: A flowchart displaying the logical structure of the experimental design of the PAST conditions, showing the number and percentage of participants that selected each response (including aggregated responses, in the grey boxes).

Results and discussion

Our hypothesis was that a deontological interpretation of the test question would be the most salient interpretation in a scenario in which Maria has already made her false statement (PAST-POSITIVE and PAST-NEGATIVE), since it is too late to *recommend* which course of action is in Maria's interest (teleological reading), but not too late to judge whether her behaviour was *improper* (deontological reading). Furthermore, the follow-up data collected in Experiment 1 indicated that participants' deontological intuitions align with the predictions non-factive accounts. On the basis of these observations, we preregistered (see Appendix) the following prediction: that the proportion of participants deeming Maria's statement permissible would be higher in the past-tense scenarios than in the present-tense one. This is indeed what we found: the proportion of non-factive answers ('agree' in PAST-POSITIVE, 'disagree' in PAST-NEGATIVE) was higher in PAST-POSITIVE (82 %) than in PRESENT-POSITIVE (42%), $\chi^2_{1,N=200} = 33.96, p < .001$, and higher in PAST-NEGATIVE (73 %) than in PRESENT-POSITIVE (42 %), $\chi^2_{1,N=200} = 21.02, p < .001$.

Beyond our preregistered predictions, we made a number of interesting findings. In both past-tense conditions, the majority of participants judged that Maria was entitled to make her false statement: 82% agreed that Maria should have said that she owns the watch in PAST-POSITIVE, and 73% disagreed with the statement that Maria should not have said it in PAST-NEGATIVE; both significantly higher than chance level (binomial test against 50%, both ps

< .001, two-tailed). That is, in both conditions, an overwhelming majority of participants chose the option that is compatible with non-factive accounts, and incompatible with factive ones (77.5% overall).

Since in Experiment 1 we found that many participants interpreted the question in an unintended way, a similar worry may arise about the results of this experiment: how many of the participants who gave a non-factive answer interpreted the main task as intended? The follow-up questions cannot provide a conclusive answer, but they certainly allow us to get a sense of how the main task was interpreted. Only 10% of the subjects who gave non-factive answers (16/155) claimed that in saying something false Maria was unintentionally violating a rule, whereas the overwhelming majority (90%, 139/155) insisted that Maria violated no rule. It is thus unlikely that these responses tracked judgements about overall permissibility (whether Maria's assertion is excusable), rather than questions about assertability (whether Maria's linguistic behaviour was appropriate) – which would undermine a non-factivist interpretation of the results.¹⁵ While follow-up questions cannot conclusively rule out every potential unwanted interpretation, this test corroborates our hypothesis that when potential sources of ambiguity are addressed, laypeople's judgements align with non-factive norms rather than factive ones.

So far, however, we have only tested people's judgements against a particular case: the ROLEX UNLUCKY vignette from Turri's (2013) study. Perhaps responses were affected by the ambiguity of the verb 'should' only in this particular experimental setting, and people would have stuck to factive responses if we had employed a different vignette. To address this worry, we set out to test people's judgements against other vignettes involving unlucky assertions.

¹⁵ One may still wonder whether phrasing the follow-up questions differently would have led to significantly different results. A referee pointed out that some of the expressions that we employed (such as "the rules governing the conversation") may be too technical to be meaningfully interpreted by participants, and that some asymmetries in the response options (in Q4, response 1 is longer than response 2) may have affected participants' responses. To address these worries, we ran an extra test on the PAST-NEGATIVE version of the task (N=150, of which 19 were excluded following the usual exclusion criteria). We rephrased the response options of the follow-ups, substituting technical expressions with plain English alternatives, and eliminating the asymmetry in length (see Appendix). Of the participants who made a non-factive judgement (76%), 74% agreed that "it was [Maria's] responsibility to answer based on the evidence available to her, so she did the right thing"; and only 26% claimed that "it was her responsibility to say the truth, but she is excusable for getting things wrong". This corroborates the results of the main experiment: most participants indicate that they are not simply excusing the violation of a factive rule, but instead that Maria is complying with a non-factive one. Similarly, of the 24% who gave factive judgements, only 16% indicated that they interpreted the 'should' deontologically, whereas 84% stated that they interpreted the 'should' teleologically. In short, we found that our results do not seem to be contingent on how the follow-up questions are phrased: the response patterns are almost identical when more ordinary terminology is employed.

4. Experiment 3

The aim of this last experiment is twofold. First, we aim to verify whether people’s preference for non-factive accounts is robust across different vignettes. Second, we aim to check whether other studies that found evidence supporting factive norms were affected by ambiguities in the test question, as we found for Turri (2013). To test both hypotheses, this experiment employs vignettes from previous studies, but modified so as to prompt participants to interpret the test question as intended. If people’s preference for factive norms remains stable in this modified setting, then we would have new and more reliable evidence in favour of factive norms of assertion. If instead we find opposite response patterns, there would be strong reasons to suspect that the preference for factive norms found by these previous studies was simply a by-product of an unintended interpretation of the research question, as suggested by the results of our previous experiments.

Given our aims, a vignette from a previous study will be suitable to be retested only if it features an unlucky assertion (followed by a should-question in the present tense). Only a few experiments meet this requirement, namely Turri’s (2014a, Exp.3) COFFEE vignette and Turri’s (2018) CAR vignette (already in Kneer, 2018).¹⁶ In the COFFEE vignette, however, it is not clear that the assertion is unlucky, because the vignette does not specify whether the protagonist is justified in forming the relevant belief, and whether the proposition is false. To see why, we report the relevant vignette below (*italics ours*):

COFFEE¹⁷

Mallory manages an independent coffee shop. One of her customers is interested in the history and culture of coffee. The customer asks Mallory whether the coffee is from Colombia. *Mallory has evidence that the coffee is from Colombia. Mallory doesn’t know whether the coffee is from Colombia.*

The sentences in italics fail to establish that it is false that the coffee is from Colombia, or that Mallory has good reasons to believe that the coffee is from Colombia. (From what is stated in the scenario, it might be *true* that the coffee is from Colombia, and Mallory might have weak or even defeated evidence that the coffee is from Colombia.)

¹⁶ If one excludes studies that feature the ROLEX UNLUCKY vignette or modified versions of it (like Turri & Park, 2018).

¹⁷ This vignette is not reported in Turri (2014a): it is merely described, and rather indirectly. We reconstructed its original formulation from personal communication with John Turri.

To make sure that Mallory's epistemic status would be interpreted as a justified true belief, we amended COFFEE so as to specify which reasons Mallory has to believe that the coffee is from Colombia, and to make it explicit that it is false that the coffee is from Colombia. To avoid a teleological interpretation of the should-question, we manipulated tense (a measure whose efficacy we tested in Experiment 2) in both the COFFEE and the CAR vignette. And to add to the variety of the repertoire, we added a third vignette, MUSIC, in which the speaker's unlucky assertion is based on reliable, albeit false, testimonial evidence (see Appendix for all three vignettes).

Participants

Data was collected until the preregistered criteria were met, resulting in 328 participants overall (64% female, age $M=33$ years). Participants received £0.20 for an estimated two minutes of their time (£6/h).

Design, materials, and procedure

The study had three between-subjects conditions (COFFEE, CAR, MUSIC), one for each vignette. Participants had to read the vignette and complete a simple task, designed after the PAST-NEGATIVE condition from Experiment 2:

[Mallory/Bill/Josie] should not have said that [the coffee is from Colombia/Jill drives an American car/ the name of the band playing in the square is Babadooks]
[Agree/Disagree]

Note that we did not test participants with the PAST-POSITIVE version of the task, since in Experiment 2 we found no significant difference in response patterns between the PAST-NEGATIVE and PAST-POSITIVE conditions. We opted for the PAST-NEGATIVE version because it is the most challenging for our account: (descriptively) fewer participants picked a non-factive option in this condition, so this method of questioning is the most likely to find evidence *against* our hypothesis.

Participants were then redirected to a new page, where they had to answer a follow-up question analogous to the one in Experiment 2,¹⁸ and then to a third page where they answered the usual control questions (adapted to the scenario).

Our hypothesis is that laypeople’s judgements align with non-factive accounts when scenario and test questions are interpreted as intended. Consequently, we predicted that the majority of participants (significantly more than 50%) would disagree with the claim that the protagonist should not have made the false assertion, also when we controlled for how the question is interpreted (i.e. also when we excluded participants who indicated that they interpreted the answer non-deontologically in the follow-up task).

Figure 4: Percentages of participants choosing the non-factive option (‘disagree’) in each scenario, before (left,) and after (right) removing participants who reported interpreting the test question non-deontologically (i.e. whose follow-up answers were incompatible with the *assertability assumption*). Error bars represent 95% confidence intervals.

Results and discussion

Our findings were consistent with our preregistered predictions (see Appendix). In every condition, the overwhelming majority of participants disagreed with the statement that the protagonist should not have made the relevant assertion (i.e. they selected the non-factive option): 92% in COFFEE, 87% in CAR, and 92% in MUSIC (all significantly higher than chance level, binomial, two-tailed, all three $ps < .001$). Overall, the overwhelming majority of participants

¹⁸ In the new follow-up, we replaced “rules of the conversation” with “conversational norms”, and we altered the tense of the conditional (“would have failed” instead of “would fail”). The text of the vignette was displayed on each new page (above the follow-up and control questions), in order for participants to be able to remember the name of the protagonist, the content of the assertion and the features of the context.

(90%) picked the option that is aligned with non-factive accounts, and incompatible with factive ones.¹⁹

These proportions are even more pronounced when we exclude the participants who indicated (in the follow-up task) that they interpreted the question non-deontologically, leaving 100 participants per condition: 96% in COFFEE, 97% in CAR, 98% in Music (binomial test against 50%, all p s < .001, two-tailed). Notice that the participants excluded in this way are a small minority (about 9%), corroborating our assumption that the past-tense design is a reliable way to probe people's intuitions about whether unlucky assertions are *assertable*, and to exclude answers based on other kinds of normativity (*teleological* normativity for factive answers, and *excusability* for non-factive ones).

Let us recapitulate. In the opening of this study, we identified a crucial limitation in the evidence collected so far in support of factive accounts of assertion: the test questions employed in these studies are ambiguous between deontological and teleological readings. Experiment 1 indicated that participants tend to interpret this sort of test question teleologically, calling into question the conclusions drawn in these studies. In experiment 2, where measures were applied to prompt a deontological reading of the test question, participants preferred the option corresponding to non-factive views. The results of Experiment 3 corroborate these findings, lending further support to the idea that previous evidence found in support of factive accounts was the by-product of an ambiguous questioning method. Our results indicate that people actually have the opposite intuition: whenever participants are prompted to interpret the test question deontologically, their judgements align with non-factive accounts.

5. General Discussion

We have accomplished three main things in this paper. First, against what has been assumed in previous studies that claimed to have found experimental support for factive accounts of assertion, we have provided evidence that judgements about what an agent 'should say' do not always track judgements of assertability (calling in to question the 'assertability assumption'). The upshot is that the conclusions about the norm of assertion drawn in these studies may well be true, but rest on shaky grounds. Second, we have introduced an alternative experimental design,

¹⁹ The results of the studies that originally employed these vignettes might be of interest for the reader. In Turri's (2018) CAR condition ("plain should" version), only 4% of participants judged that the false assertion "should be made", and in the Turri (2014) COFFEE condition only 6% agreed (to some degree) that the assertion "should be made" (but note that this latter vignette was significantly amended in our study, as specified in the main text).

arguing (and testing) that it reliably tracks intuitions about assertability. Third, we have found that once this design is adopted, laypeople overwhelmingly judge unlucky assertions to be permissible: they think that it is permissible to say something false, if you are justified in believing that what you say is true. Since we implemented measures to investigate how the test question is interpreted, our findings can be considered more reliable than the ones obtained by previous studies. Our results indicate that factive accounts of assertion, and consequently the knowledge account of assertion, are inconsistent with laypeople’s linguistic behaviour. Non-factive accounts of assertions are better suited to explaining such behaviour, in line with the findings of every study that did not employ Turri’s design (Kneer, 2018; Reuter & Brössel, 2018).

How experimental evidence about unlucky assertions bears on the existing philosophical debate on the norm of assertion, however, is a relatively complex matter. In our concluding remarks, we will attempt to clarify where our results stand with respect with the broader debate. In 5.1, we stress that adding some epicycles to the influential ‘knowledge rule’ account would not suffice to square its predictions with the evidence we collected. In 5.2 we clarify why it is unlikely that participants are simply *excusing* false assertions, rather than indicating that they are permissible. Section 5.3 addresses the worry that *outcome bias* may have distorted our results, and section 5.4 elaborates on the stability of our participants’ judgements.

5.1 Knowledge norm and factivity

Our results are especially significant for the contemporary debate on social epistemology, because they undermine previous experimental evidence in favour of the influential “knowledge account of assertion”. According to this view, assertion is governed by the ‘knowledge rule’ : “one must: assert that p only if one knows that p ”. This rule is typically taken to be *factive*. As Turri puts it, “since knowledge requires truth, knowledge is ipso facto a factive norm”. This is why unlucky assertions have traditionally been taken to be a *litmus test* for whether this view is right: if they are judged to be permissible, then the knowledge rule is (in some way or another) in trouble (Lackey, 2007, p. 603; Douven, 2006, pp. 476–7; Hill & Schechter, 2007, p. 109).

A first way to square the knowledge rule with our findings is to rethink what we mean by “knowledge”. We may postulate that knowledge does not require truth. Let us call this conception of knowledge “knowledge*”, to differentiate it from the traditional conception, that requires truth. The knowledge* rule would be compatible with our findings: unlucky assertions can be known* by the speaker even if they are false. But this revised norm would be radically different from the knowledge rule of assertion proposed in the literature – which, as Turri correctly points out, is by definition a factive norm. So construed, the knowledge-norm* would

collapse into its (non-factive) rival accounts. Some of the most prominent rivals of the knowledge rule require every ingredient for knowledge but truth; for instance, Kvanvig's (2009) justification rule requires knowledge* for permissible assertion (see also Coffman, 2014). In short, the revised knowledge* rule would only be able to accommodate the evidence by switching sides, accepting all the main contentions of the opponents of the knowledge rule: that a permissible assertion needs to be justifiably believed, but does not need to be true (nor known, in the traditional sense).

There is an alternative approach that would preserve a strong connection between knowledge and assertion: to rethink what we mean by “norm”. We may revise our hypothesis, and redefine the ‘norm of assertion’ as a norm that sets a standard for what is *optimal* to assert, rather than what is *permissible* to assert (Jackson, 2012; Mehta, 2016; Turri, 2014a)²⁰. This would offer a neat explanation for our findings: participants judge that unlucky assertions “should not” be made in the teleological sense (because unknown assertions are not *optimal*) but not in the deontological sense (because unknown assertions are nonetheless *permissible*). However, this solution has two unsurmountable problems. First, this is a radical shift of paradigm that defeats the very rationale behind the empirical studies. Proponents of non-factive norms are not willing to deny that *unlucky assertions* are suboptimal: their position is that, although suboptimal, justified but false assertions are permissible to make (Douven, 2006, pp. 476–477; Marsili, 2018, pp. 639–641; McKinnon, 2015, p. 160). Interpreting the notion of ‘norm’ in this way dissolves the very disagreement that the experimental research aims to settle. Second, this revision still leaves open the research question that was originally raised in the literature, namely the question of what makes an assertion *permissible*. Even if we were to agree that knowledge sets the standard of *optimal* assertion, the question remains open as to what sets the standard for *permissible* assertion. The findings of this study would still suggest that the question about permissibility is best solved by appealing to a non-factive rule, rather than a factive knowledge rule. Given that the research

²⁰ Although our critical target is the knowledge-rule *as it is understood in the academic debate*, some clarifications about how Turri understands this hypothesis are in order, since he authored of much of the experimental work that we discuss. Turri (2014a, 564) explores a double-standard view, according to which knowledge is only required for “good” assertion (“you well assert that Q only if you know that Q”), whereas “reasonable belief sets the standard for *permissible* assertion (that is, you may assert Q only if you reasonably believe Q)”. In subsequent work Turri (2016, 65-7) rejects this hypothesis. Crucially, his experimental studies take a more orthodox stance: in these studies the knowledge-rule is interpreted as *incompatible* with non-factive views (Turri, 2013, 281-2; 2014b, 386; 2015b, 4010; Turri & Blouw, 2015, 617), meaning that compliance with it is not merely “supererogatory”. Finally, an unrelated issue is the *function* of the norm of assertion: Turri argues the point of having a knowledge-rule in place is that this enables reliable knowledge transmission (Turri 2016a, 16-20; 2016b). For more on Turri’s interpretation of the hypothesis, see Turri (2016, 62-67).

question asked in the ongoing academic debate is the one about permissibility (Turri, 2016, p. 62; Reuter & Brössell, 2018, p. 308), the significance of our findings would remain the same. In sum, the envisaged revisions of the hypothesis would be rather cosmetic, and would not really change the significance of our findings: they mean trouble for the knowledge rule, as it is understood in the academic debate.

5.2 Excuse validation

In a recent study, Turri and Blouw (2015) found that “when an agent blamelessly breaks a rule, it significantly distorts people’s description of the agent’s conduct, [and] roughly half of people deny that a rule was broken” (Turri & Blouw, 2015) – a phenomenon that they call *excuse validation*. In light of these findings, Turri (2013:287-9) argues that laypeople may sometimes judge unlucky assertions to be permissible just because they engage in excuse validation: they confuse what is *permissible to do* with what is *impermissible but excusable to do*. If our participants engaged in excuse validation, then the evidence we found in support of non-factive accounts would be inconclusive.

We are confident that we can successfully address this worry. First, the study of excuse validation found that the proportion of people engaging in excuse validation follows some recurring patterns: about half the subjects display this distortion of judgements (between 49% and 56% in Turri & Blouw, 2015). Our study found very different proportions: 78% picked a non-factive option in Exp 2, and 90% overall in Exp 3. These response patterns simply do not match the patterns one would expect if participants engaged in excuse validation. More importantly, we anticipated this worry and controlled for excuse validation. Turri and Blouw (2015) found that excuse validation disappears when participants are allowed to judge that the excusable protagonist has “violated the norm *unintentionally*”. Consequently, in our follow-up questions we allowed participants to judge that the protagonist “violated a conversational norm, but did it inadvertently”. In both Experiment 2 and 3, almost no participant picked this option, corroborating the hypothesis that our findings were not distorted by excuse validation. In other words, our follow-up questions allowed us to differentiate between judgements about which assertions are *appropriate*, and judgements about which assertions are *inappropriate but excusable/blameless*.

These results also thwart other familiar strategies to square factive accounts with the intuitive permissibility of unlucky assertions, such as Williamson’s (2000, forth) distinction between *excusable assertions* and *permissible assertions*, or DeRose’s (2002) distinction between *primary* and *secondary* propriety. Our experimental data is hard to square with these interpretations, because

we found that unlucky assertions are judged to be permissible *simpliciter*, as opposed to impermissible but reasonable, justifiable, or otherwise excusable.

5.3 Outcome Bias

In a recent paper, Gerken (2018) challenges the significance of experimental findings supporting factive epistemic norms on different grounds. His suggestion is that these studies may have been distorted by the ubiquitous phenomenon of outcome bias. Converging evidence from several disciplines (see Gerken, 2018, §4, for a comprehensive review) has established that knowing the outcome of a given process influences a subject's evaluation of that process, even when a subject is instructed to evaluate the process *independently* of its outcome. To give an example, in their seminal study, Baron and Hershey (1988) asked their participants to evaluate the decision to operate on an elderly man, given that the operation's risk of failure was 8%. They found that "cases in which the outcome was success [...] were rated higher than matched cases in which the outcome was failure" (1988, p. 572).

Gerken points out that we should expect a similar distortion of judgements in experimental studies exploring whether a given process (e.g. making an assertion) is governed by a factive norm. In studies of factive norms, participants are typically informed of the outcome of the process that they have to evaluate (e.g. being told whether the assertion is true). For instance, in Experiment 1 participants are informed that Maria's assertion is false. Given the ubiquity of outcome bias, we should expect that knowledge of the outcome of Maria's assertion will affect participants' evaluation of Maria's action (their evaluation of whether she was entitled to make her assertion), skewing their response patterns towards a negative evaluation of the action.

While our study was not designed to test Gerken's hypothesis, our results might tell us something about its viability. If factive response patterns were *merely* the by-product of outcome bias, then participants should have deemed Maria's statement unassertable independently of whether they interpreted the main question teleologically or deontologically, since in both cases information about the outcome is available. Instead, participants only provided a negative evaluation when the test question was interpreted teleologically. We can therefore affirm that, if responses were affected by outcome bias at all, this bias had a minor effect on them. That being said, there is also a sense in which our study corroborates Gerken's suspicion that the evidence for factive norms was defective: we have, after all, found evidence suggesting that people did not interpret the test question in the intended way.

5.4 Stability

Participants changed their responses quite radically in response to small changes in the experimental setup. Our studies found near-ceiling effects for almost every manipulation involving unlucky assertions, which is the opposite of what Turri and colleagues found. This may be taken to suggest that laypeople’s judgements about unlucky assertions are extremely instable, and that employing such judgements is not a reliable method for investigating the norm of assertion.

This objection is misguided. If, as we hypothesized and checked, Turri’s experimental setup probed people’s intuitions about whether an assertion is optimal or successful, as opposed to whether it is permissible, then it is not surprising that people’s response patterns changed radically when we corrected for this defect. To put it simply: since our study asked a different question, it is not worrying (or surprising) that participants gave different answers.²¹ In addition to this, our results align with every other study (Kneer, 2018, Reuter & Brössel, 2018) that did not employ the questioning method adopted by Turri and colleagues. In fact, our study is the first to offer a plausible explanation for this difference in results: an undetected ambiguity in the test question adopted in various studies by Turri and colleagues.

That being said, our findings also provide reasons to be cautious when it comes to drawing conclusions from laypeople’s responses to surveys (for similar worries concerning experimental epistemology, see e.g. Dinges, 2016; Gerken, 2018; Nagel, 2010; Nagel, Juan & Mar, 2013). Specifically, our follow-ups indicated that the term “should” was not interpreted as researchers intended it to be interpreted, putting into question the key conclusions that were drawn from this data. Given this, one may wonder if similar problems arise for our study: participants could have similarly interpreted our tasks (including the follow-up tasks) in an unintended way.

We are not assuming that our interpretation of the results is immune from doubt. Nonetheless, we would argue that our conclusions stand on firmer ground than those drawn by the studies that we have criticized, for two main reasons. First, in our study we collected both direct evidence (the main task) and higher-order evidence (follow-up tasks) for people’s intuitions. Although neither is immune from doubt, their mutual consistency makes a sceptical approach towards our interpretations less warranted. Second, while the studies that we criticized used the same present-tense “should-judgment” task, we used a wider range of expressions and tenses, both in the main task and in the follow-up questions (cf. Experiment 2), whose variety

²¹ In fact, if participants interpreted the test question *teleologically* and indicated what is *optimal* for the speaker to do, it is also not surprising that they consistently preferred known statements over the competing epistemic standards, given that knowledge is intuitively a more desirable epistemic standard than the other ones tested in these studies (justification, belief, and truth).

was explicitly designed to minimise the risk of misunderstanding. Hence, while it is important to keep in mind that inferring epistemic norms from laypeople’s judgements is a complex and fallible endeavour, there is solid ground for assuming that our methodology significantly reduces uncertainty about participants’ interpretation of the main task.

6. Conclusion

Assertions are both ordinary and important: it is by making them that we share information, coordinate our actions, and communicate our beliefs and desires. A large body of experimental studies appears to provide converging evidence that assertions are subject to a *factive* norm: you are entitled to assert a proposition p only if p is true. All these studies, however, assume that we can treat participants’ judgements about what an agent ‘should say’ as evidence of their intuitions about assertability. Our study has found evidence that this assumption does not hold in experimental setups employing should-judgements in the present tense: participants do not interpret the test questions as questions about assertability, but rather as questions about what is in the agent’s interest to do. This casts serious doubt on the significance of the evidence gathered so far in support of factive norms of assertion (and the knowledge norm specifically). To address it, we introduced a novel experimental design to investigate laypeople’s assertability intuitions, which prompts participants to interpret the test questions in the intended way. Applying this alternative method, we found that laypeople instead favour *non-factive* accounts: they think that it is permissible to make a false statement, if you reasonably believe it to be true. Since we checked how participants interpreted the test question, the evidence gathered in our studies is a more reliable indicator of laypeople’s intuitions about the norms governing human communication. While further work in this field will be required in order to settle the disagreement about the norm of assertion, the empirical case for factive accounts of assertion now seems to be seriously undermined.

Acknowledgments

The authors would like to thank James Lewis, Ivan Milič, Manuel García-Carpintero, Vladimir Krstić, Joachim Horvath, Michele Palmira and Margherita Tattini for helpful comments on earlier versions of this draft, and John Turri for clarifying some details of the experiments we discuss in the paper. This work was supported by the German Research Foundation (DFG), project number 39130476.

Declaration of interest: none

Appendix: Supplementary Material

Link to archived Data:

<https://osf.io/v9exp/>

Links to preregistrations:

Experiment 1: https://osf.io/z395n/?view_only=e0aff96248b8430c80e37798203f959b

Experiment 2: https://osf.io/zmxp7/?view_only=c918830e05ff467e855417ed7cce6092.

Experiment 3: https://osf.io/dzp86/?view_only=b3e5d0758e754c388cf1dc518ac0918e

Vignettes and questions, Experiment 2

EXTRA CONDITION

As mentioned in footnote 12, we tested an extra condition: PRESENT-YES (N=100). This condition was identical to the PRESENT-POSITIVE condition, but the response options to Q1/Q2/Q3 were [yes/no], rather than [agree/disagree]. As expected, the transition from [yes/no] to [agree/disagree] response options did not affect participants' patterns of response to Q1: the proportion of positive answers did not differ between PRESENT-YES and PRESENT-POSITIVE, $\chi^2_1, N=200 = 0.08, p = .78$.

In the main text, we reported that the proportion of participants choosing non-factive responses was higher in past-tense conditions (PAST-POSITIVE and PAST-NEGATIVE) than in PRESENT-POSITIVE. Since the responses to Q1 in PRESENT-YES (44%) were almost identical to those of PRESENT-POSITIVE (42%), it should not come as a surprise that the proportion of participants choosing non-factive answers in PAST-POSITIVE (82 %) was also higher than in PRESENT-YES, $\chi^2_1, N=200 = 30.97, p < .001$, and similarly in PAST-NEGATIVE (73 %), $\chi^2_1, N=200 = 18.6, p < .001$.

PAST VERSION of ROLEX UNLUCKY

Maria is a watch collector. She owns so many watches that she cannot keep track of them all by memory alone. So she maintains a detailed inventory of them. She keeps the inventory up to date. Today Maria is having guests over for dinner. Soon after dinner is served, one of her guests asks, "Maria, do you have a 1990 Rolex Submariner in your watch collection?" Maria consults her inventory. It says that she does have a 1990 Rolex Submariner in her collection. Accordingly,

Maria replies: “Yes, I do have a 1990 Rolex Submariner.” But this is one of those rare cases where the inventory is wrong: she does not have one.

ALTERNATIVE FOLLOW-UP QUESTIONS (footnote 14)

Maria should not have told her guest that she has a 1990 Rolex Submariner in her collection ...

1. ...because it was her responsibility to say the truth, but it is false that she has a Rolex 1990 Submariner in her collection
2. ...because she was trying to say the truth, but it is false that she has a Rolex 1990 Submariner in her collection

It was permissible for Maria to tell her guest that she has a 1990 Rolex Submariner in her collection...

1. ...because it was her responsibility to answer based on the evidence available to her, so she did the right thing, given the circumstances
2. ...because it was her responsibility to say the truth, but she is excusable for getting things wrong, given the circumstances

Vignettes, Experiment 3

COFFEE

Mallory manages an independent coffee shop. One of her customers is interested in the history and culture of coffee.

The customer asks Mallory whether the coffee is from Colombia. Mallory checks the coffee beans label, which says that the coffee is indeed from Colombia, so she replies:

“The coffee is from Colombia”

However, Mallory does not know that the labels have been mixed up at the factory, and that the coffee beans are actually from Guatemala.

CAR

Bob has a friend, Jill, who has driven a Buick for many years. A Buick is an American car. Bob therefore thinks that Jill drives an American car.

He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill replaced it with a Mercedes, which is a German car.

One day Bob’s wife asks him, “Does your friend Jill drive an American car?”

Bob replies: “Yes, Jill drives an American car”

MUSIC

Josie and Gerald are having dinner at a restaurant. During their meal, they can hear loud music coming from the main square, just down the street. After a while, Gerald asks Josie: “Do you know who is playing this music?”

Josie checks the events section of the local newspaper, to see if any concerts are advertised. The newspaper states that tonight the band ‘Babadooks’ is playing live in the main square. Accordingly, Josie answers: “The name of the band playing in the square is Babadooks.”

However, unbeknownst to Josie, the local newspaper is wrong: the band Babadooks has cancelled at the last minute, and their concert has been replaced. A different band, ‘The Cockadoodledoots’, is playing in the main square instead.

References

- Alston, W. P. (2000). *Illocutionary acts and sentence meaning*. Cornell University Press.
- Bach, K. (2008). Applying pragmatics to epistemology. *Philosophical Issues*, (18), 68–88.
- Baron, J., & Hershey, J. C. (1988). Outcome Bias in Decision Evaluation. *Journal of Personality and Social Psychology*, 54(4), 569–579. <https://doi.org/10.1037/0022-3514.54.4.569>
- Boult, C. (2017). Epistemic normativity and the justification-excuse distinction. *Synthese*, 194(10), 4065–4081. <https://doi.org/10.1007/s11229-016-1127-8>
- Cappelen, H. (2011). Against Assertion. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 21–48). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof>
- Carter, J. A. (2017). Assertion, uniqueness and epistemic hypocrisy. *Synthese*, 194(5), 1463–1476. <https://doi.org/10.1007/s11229-015-0766-5>
- Carter, J. A., & Gordon, E. C. (2011). Norms of Assertion: The Quantity and Quality of Epistemic Support. *Philosophia*, 39(4), 615–635. <https://doi.org/10.1007/s11406-011-9317-6>
- Chomsky, N. (1965). *Aspects of the Theory of Syntax* (Vol. 16). MIT Press.
- Coffman, E. J. (2014). Lenient Accounts of Warranted Assertability. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: New Essays on Action, Belief, and Assertion*.
- DeRose, K. (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2), 167–203.
- Dinges, A. (2016). Epistemic Invariantism and Contextualist Intuitions. *Episteme*, 13(2), 219–232. <https://doi.org/10.1017/epi.2015.36>
- Douven, I. (2006). Assertion, Knowledge, and Rational Credibility. *The Philosophical Review*, 115(4), 449–485. <https://doi.org/10.1215/00318108-2006-010>
- Engel, P. (2008). In what sense is knowledge the norm of assertion? *Grazer Philosophische Studien*, 77(1), 99–113.
- García-Carpintero, M. (2004). Assertion and the Semantics of Force-Markers. In C. Bianchi (Ed.), *The Semantics/Pragmatics Distinction* (pp. 133–166).
- Gerken, M. (2012). Discursive justification and skepticism. *Synthese*, 189(2), 373–394. <https://doi.org/10.1007/s11229-012-0076-0>

- Gerken, M. (2017). *On folk epistemology: how we think and talk about knowledge*. Oxford University Press.
- Gerken, M. (2018). Truth-Sensitivity and Folk Epistemology. *Philosophy and Phenomenological Research*, 1–23. <https://doi.org/10.1111/phpr.12515>
- Goldberg, S. C. (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press.
- Golombek, H. (1977). Zugzwang. In *Golombek's encyclopedia of chess*. Crown Publisher.
- Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hawthorne, J. (2004). Knowledge and Lotteries. In *Knowledge and Lotteries* (pp. 1–49). Oxford University Press. <https://doi.org/10.1093/0199269556.001.0001>
- Hill, C. S., & Schechter, J. (2007). Hawthorne's Lottery Puzzle and the Nature of Belief. *Philosophical Issues*, 17.
- Hindriks, F. (2007). The status of the knowledge account of assertion. *Linguistics and Philosophy*, 30(3), 393–406. <https://doi.org/10.1007/s10988-007-9019-5>
- Jackson, A. (2012). Two ways to put knowledge first. *Australasian Journal of Philosophy*, 90(2), 353–369. <https://doi.org/10.1080/00048402.2011.587438>
- Kemp, G. (2007). Assertion as a Practice. In D. G. and G. Siegart (Ed.), *Truth and Speech-Acts* (pp. 106–129). Routledge.
- Kneer, M. (2018). The Norm of Assertion: Empirical Data. *Cognition*, 177(July 2017), 165–171. <https://doi.org/10.1016/j.cognition.2018.03.020>
- Kvanvig, J. L. (2009). Assertion, Knowledge, and Lotteries. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge*. Oxford University Press.
- Lackey, J. (2007). Norms of Assertion. *Noûs*, 41(4), 594–626. <https://doi.org/10.1111/j.1747-9991.2007.00065.x>
- Levinson, S. (1983). *Pragmatics*. New York: Cambridge University Press.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Clarendon Press.
- Madison, B. J. C. (2017). On justifications and excuses. *Synthese*, 1–12. <https://doi.org/10.1007/s11229-017-1418-8>
- Marsili, N. (2018). Truth and assertion: rules versus aims. *Analysis*, 78(4), 638–648.

<https://doi.org/10.1093/analys/any008>

- Marsili, N. (2019). The norm of assertion: a “constitutive” rule? *Inquiry*.
<https://doi.org/10.1080/0020174X.2019.1667868>
- McKinnon, R. (2012). What I Learned in the Lunch Room about Assertion and Practical Reasoning. *LOGOS & EPISTEME*, 4(2011), 565–569.
- McKinnon, R. (2013). The supportive reasons norm of assertion. *American Philosophical Quarterly*, 50(2), 121–135.
- McKinnon, R. (2015). *Norms of Assertion: Truth, Lies, and Warrant*. Palgrave-Macmillan.
- Mehta, N. (2016). Knowledge and Other Norms for Assertion, Action, and Belief: A Teleological Account. *Philosophy and Phenomenological Research*, 93(3), 681–705.
<https://doi.org/10.1111/phpr.12222>
- Mikhail, J. (2011). *Elements of Moral Cognition: Rawls’ Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*. Cambridge University Press.
- Musolino, J. (2009). The logical syntax of number words: Theory, acquisition and processing. *Cognition*, 111(1), 24–45. <https://doi.org/10.1016/j.cognition.2008.12.008>
- Nagel, J. (2010). Knowledge ascriptions and thinking about error. *The Philosophical Quarterly*, 60(239), 286–306. <https://doi.org/10.1111/j.1467-9213.2009.624.x>
- Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129(3), 652–661. <https://doi.org/10.1016/j.cognition.2013.02.008>
- Noveck, I. (2001). When children are more logical than adults. *Cognition*, 78, 165–188.
- Noveck, I., & Sperber, D. (2004). *Experimental pragmatics*. Springer.
- Pagin, P. (2016). Problems with Norms of Assertion. *Philosophy and Phenomenological Research*, 93(1), 178–207. <https://doi.org/10.1111/phpr.12209>
- Pelling, C. (2013). Assertion, Telling, and Epistemic Norms. *Australasian Journal of Philosophy*, (September), 1–14. <https://doi.org/10.1080/00048402.2013.798340>
- Reuter, K., & Brössel, P. (2019). No Knowledge Required. *Episteme*, 16(3), 303–321.
<https://doi.org/10.1017/epi.2018.10>
- Reynolds, S. L. (2002). Testimony, Knowledge, and Epistemic Goals. *Philosophical Studies*, 110(2), 139–161. <https://doi.org/10.1023/A:1020254327114>
- Royzman, E., & Hagan, J. P. (2017). The shadow and the tree. In J.-F. Bonnefon & B.

- Trémolière (Eds.), *Moral Inferences*. Psychology Press.
- Schechter, J. (2017). No Need for Excuses Against Knowledge-First Epistemology and the Knowledge Norm of Assertion. In J. A. Carter, E. C. Gordon, & B. W. Jarvis (Eds.), *Knowledge First: Approaches in Epistemology and Mind* (pp. 1–295). Oxford University Press.
<https://doi.org/10.1093/oso/9780198716310.001.0001>
- Schwarz, N. (2014). *Cognition and communication: Judgmental biases, research methods, and the logic of conversation*. Psychology Press.
- Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language* (Vol. 626). Cambridge university press.
- Turri, J. (2013). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, 129(2), 279–291. <https://doi.org/10.1016/j.cognition.2013.06.012>
- Turri, J. (2014a). Knowledge and suberogatory assertion. *Philosophical Studies*, 167(3), 557–567.
<https://doi.org/10.1007/s11098-013-0112-z>
- Turri, J. (2014b). Knowledge and the norm of assertion: a simple test. *Synthese*, 192(2), 385–392.
<https://doi.org/10.1007/s11229-014-0573-4>
- Turri, J. (2014c). Selfless assertions: some empirical evidence. *Synthese*, (October 2014), 1–23.
<https://doi.org/10.1007/s11229-014-0621-0>
- Turri, J. (2015). Assertion and assurance: Some empirical evidence. *Philosophy and Phenomenological Research*, 90(1), 214–222. <https://doi.org/10.1111/phpr.12160>
- Turri, J. (2016). Vision, knowledge, and assertion. *Consciousness and Cognition*, 41, 41–49.
<https://doi.org/10.1016/j.concog.2016.01.004>
- Turri, J. (2017a). Knowledge and Assertion in “Gettier” Cases. *Philosophical Psychology*, 5089(April), 1–41. <https://doi.org/10.1080/09515089.2016.1154140>
- Turri, J. (2017b). The distinctive “should” of assertability. *Philosophical Psychology*, 30(4), 477–485.
<https://doi.org/10.1080/09515089.2017.1285013>
- Turri, J. (2018). Revisiting norms of assertion. *Cognition*, 177(April), 8–11.
<https://doi.org/10.1016/j.cognition.2018.03.023>
- Turri, J. (2020). Truth, fallibility , and justification : new studies in the norms of assertion. *Synthese*.
<https://doi.org/10.1007/s11229-020-02558-7>
- Turri, J., & Blouw, P. (2015). Excuse validation: a study in rule-breaking. *Philosophical Studies*,

- 172(3), 615–634. <https://doi.org/10.1007/s11098-014-0322-z>
- Turri, J., & Buckwalter, W. (2017). Descartes's Schism , Locke' s Reunion: Completing the Pragmatic Turn in Epistemology. *American Philosophical Quarterly*, 54(1), 24–26.
- Turri, J., Friedman, O., & Keefner, A. (2017). Knowledge central: A central role for knowledge attributions in social evaluations. *Quarterly Journal of Experimental Psychology*, 70(3), 504–515. <https://doi.org/10.1080/17470218.2015.1136339>
- Turri, J., & Park, Y. (2018). Knowledge and Assertion in Korean. *Cognitive Science*, 1–21. <https://doi.org/10.1111/cogs.12621>
- Weiner, M. (2005). Must We Know What We Say? *The Philosophical Review*, 114(2), 227–251. <https://doi.org/10.1215/00318108-114-2-227>
- Whiting, D. (2012). Stick to the Facts: On the Norms of Assertion. *Erkenntnis*, 78(4), 847–867. <https://doi.org/10.1007/s10670-012-9383-6>
- Wiegmann, A., Samland, J., & Waldmann, M. R. (2016). Lying despite telling the truth. *Cognition*, 150, 37–42. <https://doi.org/10.1016/j.cognition.2016.01.017>
- Williamson, T. (n.d.). Justifications, Excuses, and Sceptical Scenarios. In J. Dutant & D. Dohrn (Eds.), *New Evil Demon*. Oxford: Oxford University Press. <https://doi.org/10.1017/CBO9781107415324.004>
- Williamson, T. (1996). Knowing and Asserting. *The Philosophical Review*, 105(4), 489–523.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press. <https://doi.org/10.5040/9781474284974.00620>