NO KNOWLEDGE REQUIRED

KEVIN REUTER AND PETER BRÖSSEL

kevin.reuter@philo.unibe.ch; peter.broessel@ruhr-uni-bochum.de

ABSTRACT

Assertions are the centre of gravity in social epistemology. They are the vehicles we use to exchange information within scientific groups and society as a whole. It is therefore essential to determine under which conditions we are permitted to make an assertion. In this paper we argue and provide empirical evidence for the view that the norm of assertion is justified belief: truth or even knowledge are not required. Our results challenge the knowledge account advocated by, e.g. Williamson (1996), in general, and more specifically, put into question several studies conducted by Turri (2013, 2016) that support a knowledge norm of assertion. Instead, the justified belief account championed by, e.g. Douven (2006), seems to prevail.

1. INTRODUCTION

Many of our speech acts serve the purpose of informing others about what is the case or stating what is true. These speech acts are typically called *assertions*. As such, assertions have played a central role in social epistemology. One of the most fundamental topics within social epistemology concerns the question about when people are justified to adopt a belief in light of testimonial evidence; or, more plainly, when is it admissible to trust the assertions of others? It is essential to note that this question cannot be answered without first providing an answer to the question of what the norm of assertion is. In particular, before we can expound on when we are justified to adopt beliefs on the basis of the assertions of others, it is essential to say under which conditions we are permitted to make an assertion in the first place. If the accepted norm of assertion were that one is permitted to assert p whenever asserting p would make others laugh, then we certainly should adopt strict norms concerning when it is epistemically admissible to trust the assertions of others. If, on the other hand, the norm of assertion were that one is permitted to assert p only if one knows p, then we should be very liberal concerning our trust in the assertion of others, i.e. then we might say that one is justified in believing that p, whenever someone asserts p, unless one has good reasons to doubt that a person respects that norm. If the members of the society really comply with such a strong norm of assertion, we might say that it is admissible to believe p given testimony that p even if we have counter-evidence that supports non-p. It is indeed rarely recognized that the question about when people are justified to trust the assertions of others cannot be answered without first detailing the norms of assertions.

Various accounts regarding the norms of assertion have been proposed during the last 15 years. Our investigation in this paper concerns the three, arguably, most influential accounts that have been discussed in the literature:

> (Knowledge) Assert p only if you know that p.
>
> (True Belief) Assert p only if you believe p and p is true.
>
> (Justified Belief) Assert p only if you believe p with justification.

At least since Gettier (1963) the classical account of knowledge, i.e. that knowledge is justified true belief, is considered to be problematic. Nevertheless, most, if not all advocates of (Knowledge) agree that knowing p entails having a justified, true belief in p. Given such an understanding of knowledge, all three proposals require that a person believes what she asserts. We will presuppose this as a minimal norm of assertion in this paper and disregard norms of assertion that do not imply it (but see Lackey (2007) for an account that does not require belief). These three accounts therefore differ with respect to the role they ascribe to truth and justification. (Knowledge) and (True Belief) both posit that p should not be asserted if p is false. (Knowledge) and (Justified Belief) both claim that p should not be asserted if p is not believed with justification.

The question of what the norms of assertions are in a given society is partly a conceptual and partly an empirical question, and depends on which speech acts we are willing to pin the conceptual label "assertion" on.[1] There is some debate in the literature on whether the suggested norms of assertion are constitutive for assertions or merely regulative. If they were constitutive for assertions, then it would be a conceptual matter that we do not pin the label "assertion" on speech acts for which we know that they violate the respective norm. Alternatively, these norms of assertions might only be regulative and, thus, presuppose "that there is something that counts as asserting, and [the regulative norms] tell us what an asserter ought to be aiming at when performing the speech act" (Maitra 2011: 282). However, whatever stand one takes on this issue, almost everyone agrees that these norms, if correct, have observable consequences concerning our social practice. It is for this reason that hypotheses about norms of assertion can be confirmed or disconfirmed by empirical data. For example, Williamson (1996), a champion of (Knowledge), openly admits that ordinary speakers follow (Knowledge) implicitly and that therefore empirical investigations into the practice of assertions among ordinary speakers can only reveal that (Knowledge) is the best explanation for their behavior. Similarly open to empirical research, Douven (2006), a defender of (Justified Belief), writes:

> In my view the project of determining which rule governs the practice of assertion is best conceived not as an a priori investigation into the nature of assertion but, rather, as an empirical project and that therefore any proposal made in the course of this project is subject to the exact same standards of evaluation as are employed in the empirical sciences generally. That is to say, any hypothesis stating that a particular rule governs the practice of assertion must face the linguistic data about that practice. (Douven 2006: 450)

However, even if it is universally acknowledged that discovering the norms of assertion is partly an empirical project, the empirical means that we ought to employ in this project

---

1 Apart from the descriptive question, there is, of course, also the normative question of what the norms of assertion *should* be in a given society. Although we remain largely silent on this matter, we do believe that the epistemic goals of a society are likely to provide important boundary conditions on what the norms of assertion are; see also the General Discussion.

are disputed. Most philosophers engaging in the debate rely on intuitions and linguistic data. However, philosophers usually do not collect linguistic data systematically, but often seem to cherry-pick data that supports their particular theory best. In addition, linguistic data can be accounted for in various ways, and it is not always clear whether a given norm of assertion is the best explanation for the linguistic data. Williamson (2000), for example, is well aware that many of his linguistic data can be explained via various other Gricean norms of conversation. More importantly, that these data points are not decisive in favoring one or the other theory is convincingly shown by Douven (2006). He argues that (Justified Belief) is the better, or at least equally suited account to explain the same data Williamson relies upon.

For example, one of Williamson's reasons in support of (Knowledge) consists in the following linguistic datum: the question "How do you know?" is a proper response to an assertion. According to Williamson, the linguistic data suggest that we presuppose that the asserter knows that p whenever she asserts that p. For Williamson, the best explanation for this linguistic datum is that one is permitted to assert that p only if one knows that p. Accordingly, when we ask "How do you know?" we question whether the asserter fulfills all the relevant requirements for being permitted to assert that p. Douven (2006), who argues in favor of a version of (Justified Belief), also believes that we have "to face the linguistic data" (Douven 2006: 45), but he does believe (Justified Belief) can explain the data equally well. Note, the adequate response by the asserter in response to the question "How do you know?" is to provide her reasons for why she believes p, nothing in addition is required from the asserter. Now, Douven thinks that asking "How do you know?" is simply the way we ask for reasons as to why one believes p; and we ask the way we do because typically when someone is justified in believing p, she knows that p or, at least, takes herself to know that p. Thus, according to Douven (2006: 469–71), we can account for the linguistic data by assuming that (Justified Belief) holds for assertions and by also assuming that typically, when someone is justified in believing p, this person typically knows that p (see also McKinnon 2015; Gerken 2017). It seems then that the truth about norms of assertion is underdetermined given the linguistic data that are available (see Pagin (2016) for an extended discussion on the question of underdetermination).

Given the standstill in the debate surrounding the available linguistic data, some philosophers have tried to tackle the empirical project with more potent means, the means of experimental philosophy. This turn has been called for in the literature. For example, in the SEP entry on assertion, Pagin (2015) states that

> Virtually throughout the discussion, authors have simply appealed to their own intuitions. However, this is an area where experimental philosophy would be highly relevant, since the question largely seems to concern what norms ordinary asserters accept. (Pagin, 2015)

Recently, this call for more experimental work has been followed. In particular, Turri has taken an experimental approach to answering the question of what the norm of assertion is. Such an approach is warranted in his view "because competing theories about the norm of assertion generate testable predictions" (Turri 2016: 11). Indeed, Turri has conducted most experimental work on this subject matter, in part in collaboration with several other philosophers. In a series of articles, Turri (2013, 2015, 2017) has collected empirical data that seem to support the knowledge account of assertion. Most of the

experimental work conducted for the present paper takes its inspiration from Turri's approach to tackle these questions empirically and aims to improve upon it in certain respects. Thus, in Section 2, we outline Turri's central experimental studies and criticize key aspects of them. In Section 3, we test our claims by manipulating these key aspects in Study 1. In Section 4, we investigate the relative importance of truth and justification in two further experiments. Section 5 summarizes our results and discusses possible limitations to our account.

2. PREVIOUS STUDIES AND CRITICISM

Turri himself raises what he considers to be the most pressing objection against the knowledge account of assertion.[2]

> Critics of the knowledge account report having the intuition that reasonable false assertions are perfectly fine. They claim that this intuition is "obvious" and reflects ordinary practice (Hill & Schechter 2007, 109; Douven 2006, 476ff). Stronger yet, some claim that there is "no intuitive sense" in which a reasonable false assertion is improper, and that "there is no practice" of counting them as inappropriate (Douven 2006, 480; Hill & Schechter 2007, 109). (Turri 2013: 39–40)

In order to counter this objection, Turri (2013) designed a vignette to probe whether laypeople indeed think it permissible to assert a false proposition if it is reasonable to hold that proposition.

> 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/And this is just another case where the inventory is exactly right: she does have one].

Directly following the vignette, participants answered various test questions and some comprehension questions. The comprehension questions were asked to check whether the participants had understood the vignette adequately and also to draw their attention to the relevant aspect of the vignette: the truth of a given statement. This set up was chosen to check whether the truth of an assertion had an effect on whether the participants considered the assertion admissible or not. Turri designed four experiments based on this vignette; and in all four experiments participants were asked whether Maria should say that there is a 1990 Rolex Submariner in her watch collection.[3]

---

2 No doubt, Turri has conducted several more studies in various other papers that he takes to provide additional support for the knowledge account. However, as far as we know, none of the other studies put the justified belief account to test in a similarly direct fashion.

3 To be more precise, in one of these four experiments they were asked "How should Maria answer her neighbor's question?" Participants responded on a 7-point Likert scale. The options varied whether and

The responses painted a clear picture according to which the majority of those participants, who were presented with the condition in which the inventory is wrong, claimed that Maria should not assert that she has a 1990 Rolex Submariner in her collection. Thus, people did not seem to think that Maria should assert that *p* despite being justified in believing that *p*. Only in the condition in which the inventory is correct, did people judge that Maria should say that she is in possession of the Rolex watch. On these grounds, Turri excludes (Justified Belief) as the appropriate norm of assertion (2014: 289). Given that (Knowledge) can explain various linguistic data that (True Belief) certainly cannot explain, Turri concludes that (Knowledge) must be the correct account of our norm of assertion (Turri 2016).

We believe that accepting this conclusion would be premature. In the following, we first set out our twofold criticism of Turri's experimental setup.[4] One criticism is directed at the vignette, another criticism is directed at the way people are prompted to respond to the vignette that is offered in Turri's experimental setup. These criticisms are driven by philosophical considerations alone. However, if our points of criticism are correct, then they will point at some crucial variations of Turri's setup that one has to empirically investigate.

Our first criticism concerns the vignette itself. In the vignette the participants are told that "Maria knows that the inventory is not perfect, but [that] it is extremely accurate." Here the vignette emphasizes that Maria is aware of one relevant error possibility, i.e. that the inventory is malfunctioning on the given occasion, and that she has not taken any efforts to rule it out. Thus, some participants might reason that she is not really justified (or at least not to a sufficient degree) in believing that there is a 1990 Rolex in her watch collection because she cannot exclude all relevant alternatives. Thus, our first proposal is to delete the respective sentence from the vignette. It should read:

> 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. [But this is one of those rare cases where the inventory is wrong: she does not have one/And this is just another case where the inventory is exactly right: she does have one].

Note, if (Knowledge) were the correct norm of assertion then people's responses to the revised vignette should paint a similar picture as the responses in Turri's original study. In one case, Maria does not know that she has a 1990 Rolex in her watch collection, because the relevant statement is false. If one wants to criticize this vignette at all, then one should do so for not specifying to a sufficient degree that Maria is justified in believing

---

how strongly Maria should claim or deny that the watch is in her collection. In conditions where the inventory says that Maria does have the watch (positive conditions), participants chose from the following options, left-to-right: "I definitely do have one"; "I do have one"; "I probably do have one"; "It's unclear"; "I probably don't have one"; "I don't have one"; "I definitely don't have one".

4   Pagin (2015) raises doubts about the consistency of Turri's results. In this paper we go further in putting forward a two-fold criticism against the experimental setup.

that she owns a 1990 Rolex Submariner. The vignette now does not explicitly state the inventory is extremely accurate. Thus, so the possible criticism goes, the participants might judge that Maria does not know that she possesses such a Rolex independently of the truth of the given claim. However, first, we think that the lines "So she maintains a detailed inventory of them. She keeps the inventory up to date" make it pretty clear that Maria does not need to worry about the inventory's fallibility. Second, we will see in Section 4 that this is not what is going on.[5]

Our second criticism concerns the possible response options provided in Turri's experimental setup.[6] Remember, in three of four experiments participants were asked whether Maria *should* say that there is a 1990 Rolex Submariner in her watch collection and in the fourth experiment they were asked: "How should Maria answer her neighbor's question?" We think it is improper to ask what Maria *should* say.[7]

First, asking whether Maria should say that there is a 1990 Rolex Submariner in her collection is problematic. In particular, when the participants' response is that Maria should not say that there is 1990 Rolex in her collection, then we still do not know why she *should* not say it. The reason is that the term *should* can be used to express an obligation but also to express instrumental expediency. Turri interprets *should* as expressing an obligation and he therefore thinks that the responses by the participants support (Knowledge). He writes "researchers working on the norm of assertion have accepted, often implicitly, several assumptions. [...] They assume that the rule is deontological because the 'should' expresses the concept of permission; you have permission or authority to assert only under certain conditions, and to do otherwise is impermissible" (Turri 2016: 61). Thus, the explanation that Turri favors is that Maria should not assert that there is a 1990 Rolex Submariner in her collection, because she is not permitted to do so according to the knowledge norm of assertion.

An alternative explanation is that Maria should not assert that there is a 1990 Rolex Submariner in her collection, because this would neither be expedient to the typical goals of a conversation nor expedient to her particular goals as given in the vignette. In particular, from the vignette the participant understands that Maria honestly aims at saying something true about her watch collection and that, given that the statement is false, she would not reach this goal. This reasoning, however, is independent of the norm of assertion. It might perfectly well be the case that the norm of assertion allows her to assert

---

5   In his 2017 paper, Turri replicated his findings using the same problematic vignette as stated above. He also conducted a second experiment yielding similar results in which no error possibility was explicitly mentioned. However, in that second experiment, a vignette was used in which a protagonist jumps to the premature conclusion that a shot was fired from hearing a loud bang. Hence, one could argue that the vignette describes a case in which a person is not sufficiently justified to make the claim in question. Independently of the question of whether or not this second vignette is problematic, we will see that our variation of Turri's original vignette leads to severe problems for the knowledge account.

6   We would like to thank Anna-Maria Asunta Eder for suggesting this objection and for discussions on this point.

7   We would like to add that Turri does not defend the knowledge account in the form stated above: Assert *p* only if you know that *p*. Rather, he advocates the knowledge account in the form "you should assert Q only if you know Q" (Turri 2013: 281). It is therefore not surprising that Turri has asked his participants a *should* question. It is an open question whether the knowledge account in its more general form can be properly cast in the way Turri suggests. In any case, our two-pronged criticism below applies independently of the exact framing of the knowledge account.

something false, but she nevertheless should not do so because it would not be expedient to her goals in the given situation.[8]

Second, note, even if *should* were expressing an obligation and (Knowledge) were the correct norm of assertion, from knowing p it does not follow that one should assert p. Instead the converse is true: if (Knowledge) is the correct norm of assertion then one should or ought to know that p whenever one asserts that p. This feature of (Knowledge) is used to explain why the question "How do you know?" is a proper response to an assertion: it is a proper response because if one asserts that p, then one can presuppose that the asserter knows that p, because that is what she should do in order to satisfy (Knowledge). We therefore suggest that one asks whether Maria is *permitted* to say that there is a 1990 Rolex Submariner in her watch collection. Note again, if (Knowledge) were the correct norm of assertion then the responses to the revised question should yield similar answers in response to Turri's original question. After all the following is correct: If Maria knows that there is a 1990 Rolex in her collection then she is permitted to assert it, but Maria would not be permitted to make that assertion if she does not know it. In those instances of the vignette in which the claim is false, she does not know it and therefore participants can be expected to answer that she is not permitted to say that there is such a watch in her collection.

We have raised two points of criticism against Turri's experimental setup and we argued that if (Knowledge) is the correct norm of assertion then these changes should not make a difference. First, we suggested that we should delete the sentence "Maria knows that the inventory is not perfect, but [that] it is extremely accurate." Thus the vignette no longer emphasizes that Maria is aware of one relevant error possibility. This does not change the overall setup, because Maria still does not know that there is a 1990 Rolex in her collection, because that assertion is false. Second, we suggested that we should ask whether Maria is permitted to assert that there is such a Rolex in her collection instead of asking whether she should say it. Again, from the perspective of (Knowledge), this variation is not supposed to make a difference. According to (Knowledge), if Maria does not know that there is a 1990 Rolex in her collection, then she is not permitted to make that assertion. In the following sections, we will now report our findings.

---

8   For illustration consider the false-belief task (Perner *et al.* 1987). The participants (typically children or infants) are shown two containers, and then one person places an object in one of these containers and subsequently leaves the room. A second person enters the room and switches the object into the other container and also leaves. When the first person returns, the participant is asked where will the person look for the hidden object? The participant fails the test if he or she thinks the person will look for the object at its actual place, indicating that he/she fails to understand that the person has the false belief that the object still is in the container in which it was initially hidden. Now suppose we ask the participant: should the person look for the object in the container in which it was originally placed? Here the correct answer seems to be 'No'. She should not look for the object in its original hiding spot. The explanation for that answer, however, has nothing to do with any norms of rational action or belief. She should not look at the original hiding spot simply because she will not find the object there (for a similar argument see Ben Yami *et al.* 2015). Asking for where the person should look for the object in the false belief task aims at determining what action will lead the person to reach her goal. Applied to our case at hand, when we are asking whether she should assert p if she has good reasons to belief p even if it is false, then the answer might well be that she should not assert p because contrary to her intentions she would be saying something false, but that she is of course permitted to assert p since she is justified to say so.

3. STUDY 1: EVIDENCE AGAINST (TRUE BELIEF) AND (KNOWLEDGE)

In the previous section, we discussed experimental studies that seem to provide evidence for the knowledge account of assertion. However, we also raised several objections against the empirical work, specifically about the content of the questionnaires that were designed to show the falsity of the justified belief account of assertion. We will now present new empirical work that questions the conclusions drawn from previous experiments. Given the concerns raised above, the purpose of Study 1 was twofold. First, we intended to replicate the results by Turri (2013) according to which people deny that a person may assert a false claim under certain conditions. Second, by manipulating two aspects of the vignette that we identified as crucial for these tests, we investigated whether the true belief and knowledge account indeed provide an empirically adequate theory of the conditions for assertion.

3.1 Methods

For Study 1, 213 participants (98 women, M_{age} = 33.93, SD = 10.75) were recruited on Amazon Mechanical Turk and paid a small fee for their participation. All participants were native English speakers. In order to replicate Turri's results we used a vignette almost identical to the one he used in his 2013 paper.

> Maria is a watch collector. She owns so many watches that she uses an electronic data-base to keep track of all her watches. Maria knows that the database does not function perfectly, but 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 electronic database, which states that she does have a 1990 Rolex Submariner in her watch collection. So Maria says: "Yes, I do have a 1990 Rolex Submariner in my collection." However, this is one of those rare cases in which the electronic database was malfunctioning and wrongly stated that she has a Rolex Submariner in her watch collection. Should Maria tell her guest that she has a 1990 Rolex Submariner in her collection?

The participants were presented with three options: (1) Yes. (2) No. (3) Other/Don't know. We then manipulated two aspects of the vignette. First, instead of asking the participants whether Maria should (N = 54) tell her guest that she has a 1990 Rolex Submariner in her collection, we asked them (N = 65):

> Was Maria permitted to tell her guest that she has a 1990 Rolex Submariner in her collection?

Second, we eliminated from the vignette the information "Maria knows that the database does not function perfectly, but is extremely accurate." As we have argued above, knowing that a system is fallible might affect people's responses regarding the assertion conditions of a false statement (for a related discussion on salient alternatives see, e.g. Nagel *et al.* 2013). Thus, the alternative scenario read:

> Maria is a watch collector. She owns so many watches that she uses an electronic database to keep track of all her watches. 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 electronic database, which states that she does have a 1990 Rolex Submariner in her watch collection. So Maria says: "Yes, I do have a 1990 Rolex Submariner in my collection." However, unbeknownst to Maria, the electronic database was malfunctioning and wrongly stated that she has a Rolex Submariner in her watch collection.

Again, we asked people either the *should* question (N = 56) or the *permission* question (N = 49), and presented them with three options: (1) Yes. (2) No. (3) Other/Don't know. After the main question was raised, we also asked each subject to briefly explain why they had responded the way they did. All participants were randomly assigned to either one of the four conditions of the 2 × 2 experimental design.

3.2 Results

Pearson's χ^2-tests were used to determine significant relationships between two scenarios in the 2 × 2 experimental design, operating with a 2 × 2 × 3 contingency table. Repeating Pearson's χ^2-test for various pairwise comparisons made it necessary to adjust the level of significance. Having limited our analysis to 3 pairwise comparisons, the level of significance reduced to 0.016 (using conservative Bonferroni correction). We started our analysis with the original scenarios in which Maria is aware of the fallibility of the electronic database. As can been seen in Figure 1, on the left hand side (dark grey bars), when the *should* question was asked, a majority of the participants did not believe that Maria should tell her guest that she has the Rolex watch in her collection, replicating a crucial finding of Turri (2013). However, when the question was changed from *should* (dark grey bars) to *permission* (light grey bars), a significant impact was found on people's answers. Whereas the 'Yes' responses increased from 18.5% to 76.9%, the 'No' responses dropped from 68.5% to 13.8%: (χ^2 = 43.14, p < 0.001, Cramer's V = 0.6).[9] We then compared (Figure 1, right hand side) the original scenario (doubt) with the case in which information about the fallibility of the database was deleted (no doubt). In both scenarios the *should* question was asked. The analysis revealed that changing information about the fallibility of the database also changed the response patterns significantly: the amount of 'Yes' responses rose from 18.5% to 62.5%, whereas the percentage of 'No' responses fell from 68.5% to 35.7%, (χ^2 = 23.43, p < 0.001, Cramer's V = 0.5). The third comparison contrasted the original case (should, doubt) with the scenario in which both factors were manipulated (permission, no doubt). An even stronger effect was noted in which 'Yes' responses rose to 87.8% and 'No' responses dropped to 8.2% (χ^2 = 49.76, p < 0.001, Cramer's V = 0.7, not depicted in Figure 1).

3.3 Discussion

In Study 1, we successfully replicated Turri's results, suggesting that under certain conditions, people do not believe that a person should make a false statement even if that person has reasons to believe her statement to be true. However, the data also demonstrate that

---

9   For similar results, cf. Kneer (Forthcoming), as well as Marsili and Wiegmann (Ms).
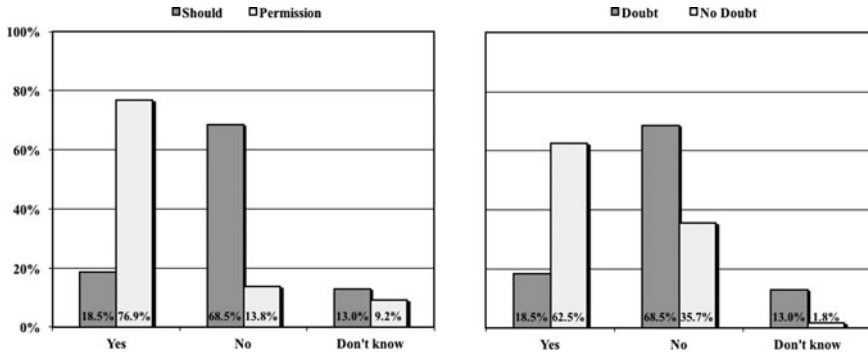
Fig. 1. Responses in per cent to the scenarios of Study 1 in which both the question type (left hand) and the scenario type (right hand) were manipulated.

this result cannot be drawn upon to support (Knowledge) or (True Belief). Indeed, our results put serious pressure on both accounts. Two factors were identified in the original experiment that might have skewed the results in Turri's favor. First, in the original vignette the protagonist knew about the fallibility of her source. In order to show that people endorse the knowledge account of assertion, such doubts should be excluded from the vignette. Second, asking whether a person *should* make a statement is a much stronger request than asking whether a person is *permitted* to make that statement – we have argued in Section 2 that an empirical inquiry into the norm of assertion should use the latter question. As hypothesized, both factors had a crucial impact on people's views. Once any belief about the fallibility of the source of Maria's knowledge was removed, a significant and substantial majority claimed that the protagonist Maria was in the right position to make a false assertion. This was true regardless of whether the participants in our study responded to the *should* or the *permission* question. When people were drawn to the permission scenario, an even greater majority (87.8%) approved of Maria's assertion, indicating that the type of request makes an important difference even if doubts about the fallibility of the source are removed.

When we kept the knowledge of the source's fallibility inside the vignette, the response patterns differed widely between the *should* and the *permission* question. Whereas the should responses matched those of Turri's original investigation, the permission responses delivered a very different verdict: a clear majority of 76.9% stated that Maria was permitted to make a false assertion, suggesting that the predictions of the knowledge account of assertion are false.

While our reasoning about the effects of changing the wording of the question and the awareness of the fallibility seems plausible, how do we know that our account is explanatorily adequate? As stated above, we also asked the participants of our study to briefly explain their responses. Importantly, when we contrasted the responses of the two scenarios in which we asked the *should* question but manipulated Maria's awareness about the fallibility of the electronic database, then the explanations for the 'No' responses reflected this difference. For example, in the original scenario, participants stated that "If she can't verify for certain then she should say 'no' until she does know for certain", "If she knows that the database is malfunctioning, then she knows that she does not own said watch thus, she shouldn't claim to", "She knew that her database malfunctions albeit

rarely and therefore she knew that she couldn't say for sure without checking her collection", "If Maria knows the database is not 100% reliable, she should not make a definitive statement about her inventory based on it." Thus, when evaluating whether or not Maria should make that assertion, readers of the vignette did not focus on Maria's knowledge that the database is highly reliable, but instead attended to the fact that Maria knows that the database is not perfect.[10] In contrast, in the scenario in which Maria's awareness of the possible failure of the database is eliminated, the majority of the participants explained their 'Yes' responses by claiming, for example, that "Since she doesn't know it is malfunctioning, then she has no reason to distrust what the machine is telling her", "She doesn't know that the system is malfunctioning, so as far as she knows, she does have that watch", "The database told her that she had one. She did not know that it was wrong so I feel she should say yes."

Interestingly, in the scenario in which Maria is aware of the malfunctioning database but asked the *permission* question, most participants are far more forgiving compared to the same scenario followed by the *should* question. Typical responses were: "She is allowed based on the information that she had at her disposal", "She had no way of knowing when the program would malfunction, and the story stated that this was a rare case", "Since her database functions correctly most of the time, it was very likely that she had the specific watch." The responses hence reveal that the bar for being permitted to make an assertion is substantially lowered. The results of Study 1 therefore seem to show that the justified belief account has not been shown to be mistaken.

4. STUDY 2: EVIDENCE FOR (JUSTIFIED BELIEF)

Advocates of (Knowledge) and (True Belief) claim that we are not permitted to assert *p* if *p* is false – even if we are justified in believing p.[11] Some of Turri's empirical studies suggested that they could well be right. However, the results of Study 1 not only demonstrate that the interpretation of previous studies was premature, they show that there are at least some cases, in which people reject the idea that truth is necessary for making an assertion when a person is clearly justified. But while the results of Study 1 do put into doubt (Knowledge) and (True Belief), they do not exclude the possibility that the justified belief account is also false. The data we obtained in Study 1 are, of course, consistent with the idea that the norm of assertion is *justified belief*, but in order to empirically investigate whether the justified belief account prevails, further experiments are required. The purpose of Study 2 is to test more generally which of the three accounts laid out above best accords with laypeople's judgements. We therefore designed two sets of scenarios that directly pitted truth against justification. More specifically, we developed and manipulated scenarios that allowed the pairwise comparison of the importance of *truth* and *justification* for the norm of assertion.

---

10   This result fits the predictions of the focal bias account of Gerken (2013).

11   Williamson acknowledges the widespread use of false morally permissible assertions but maintains that they violate the norm of assertion.

4.1 Study 2A

4.1.1 Methods

We randomly assigned 91 participants to two different scenarios.[12] Four participants were excluded for not having finished the survey. Scenario A1 (N = 44) was designed such that the main protagonist of the scenario was justified in believing p, although p was false. In contrast, in scenario A2 (N = 43), the main protagonist made an assertion that was true but which she was not justified in believing. If the justified belief account is correct, then we expect the participants to (a) approve S's making of the false assertion p when S is justified in believing p and (b) reject S's making of the true assertion q when S is not justified in believing q.

Each participant was presented with only a single vignette. After each scenario, we asked the participants whether the protagonist of the scenario was permitted to make the respective assertion.

Scenario A1: (Justified but not true)

Maria is a watch collector. She keeps all her watches in a safe and enjoys looking at her collection regularly. One day, her friend John asks her, whether she has a 1990 Rolex Submariner in her safe and, if so, could show it to him. Maria answers: "There is a 1990 Rolex Submariner in my safe." She does say so, because she had purchased that watch a few years ago and has seen that watch lately in her safe. When Maria opens the safe a few minutes later, she finds out that a burglar has stolen several watches, among them the 1990 Rolex Submariner.

Q: Was Maria permitted to tell John that there is a 1990 Rolex Submariner in her safe?
R1: Yes, because she had purchased that watch and has seen it lately in her safe.
R2: No, because there is no 1990 Rolex Submariner in her safe.
R3: Other/Don't know.

Scenario A2: (True but not justified)

Maria is a watch collector. She keeps all her watches in her safe and enjoys looking at her collection regularly. One day, her friend John asks her, whether she has a 1990 Rolex Submariner in her safe and, if so, could show it to him. Maria answers: "No, there is no 1990 Rolex Submariner in my safe." She does say so, although Maria purchased that watch a few years ago and has seen that watch lately in her safe. Maria just simply forgot this when John asked her. When Maria opens the safe a few minutes later, she finds out that a burglar has stolen several watches, among them the 1990 Rolex Submariner.

Q: Was Maria permitted to tell John that there is no 1990 Rolex Submariner in her safe?
R1: Yes, because there is no 1990 Rolex Submariner in her safe.
R2: No, because she had purchased that watch and has seen it lately in her safe.
R3: Other/Don't know.

After having been presented with the vignette and the test question, we also asked the participants two control questions: (i) Was Maria's answer true? (ii) Was Maria's answer justified?

---

12 Participants were recruited through Amazon's Mechanical Turk. The same restrictions were applied in all experiments of Study 2 as in Study 1.

4.1.2 Results

When the participants were asked to judge whether Maria was permitted to make an assertion that was false but justified (scenario A1), a clear majority of the participants answered affirmatively. 90.9% (N = 40) responded with 'Yes' while only a single person answered 'No'; three participants (6.8%) selected the 'Don't know' option. In contrast, those participants who were assigned to scenario A2 in which the protagonist Maria makes an assertion that is true but unjustified, merely 11 subjects (25.6%) thought that Maria was permitted to make the assertion. 21 (48.8%) participants responded with 'No' and a further 11 (25.6%) participants opted for 'Don't know'. The percentages are also displayed in Figure 2 (left hand side). A χ^2-test revealed that the difference in responses between the two scenarios was highly significant: $\chi^2 = 39.24$, $p < 0.001$, Cramer's V = 0.7. We also analyzed the responses of those participants who correctly responded to both control questions (N = 37). Percentages are depicted in Figure 2 (right hand side). Fisher's Exact Test was used because no subject who passed both controls answered 'Don't know': $\chi^2 = 26.93$, $p < 0.001$. In both scenarios, the responses were significantly different from chance level.

4.2 Study 2B

Study 2A strongly suggests that when truth and justification are pitted against each other, subjects judge that a person is permitted to make an assertion when she is justified and not permitted to do so when not having justification for her belief. In contrast, (True belief) and (Knowledge) make wrong predictions about the outcome of the experiment. Before we discuss these results in detail below (see Section 4.3), we would like to respond to the possible objection that the results of Study 2A might be due to some artifact of the story of the vignette or its specific wording that we used. We will respond to this objection empirically, i.e. we designed two more vignettes in which the respective importance of truth and justification was tested.

4.2.1 Methods

Eighty-six participants were recruited through Amazon Mechanical Turk and assigned randomly to one of the two scenarios stated below. Three subjects had to be excluded
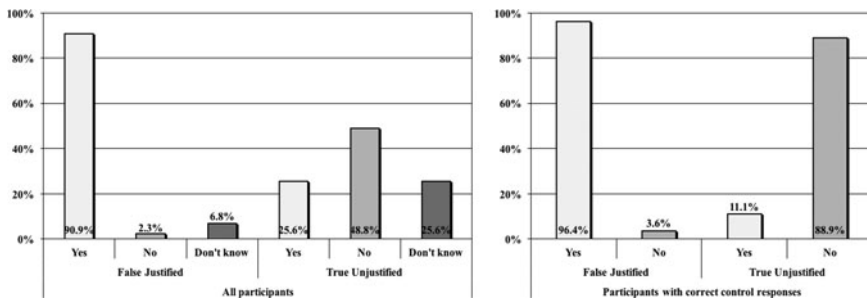

Fig. 2. Percentages of the responses for Scenario A1 (False Justified) and Scenario A2 (True Unjustified) when all participants' responses were evaluated (left hand side) and when only those responses from participants were taken into account that passed both comprehension checks (right hand side).

as they failed to complete the survey. Similarly, to Study 2A, scenario B1 was designed as to present a story in which the main protagonist Robert makes a claim for which he is properly justified but which turns out to be false. Scenario B2 was only changed in a way such that Robert is not justified in believing the assertion he makes but in which the claim turns out to be true.

Scenario B1: (Justified but not true)
  Anne and Robert go to a costume party late at night. On their way to the party, Anne asks Robert whether any of his friends are also at the party. Robert answers: "Jill is at the party". He does say so, because Jill had told Robert a few hours before that she had already selected a nice dress and was planning to come. When they arrive at the party, it turns out that Jill had changed her plans, and actually did not go to the costume party.

  Q: Was Robert permitted to tell Anne that Jill would be at the party?
  R1: Yes, because Jill had told Robert that she would go.
  R2: No, because Jill did not go to the party.
  R3: Other/Don't know.

Scenario B2: (True but not justified)
  Anne and Robert go to a costume party late at night. On their way to the party, Anne asks Robert whether any of his friends are also at the party. Robert answers: "Jill is at the party". He does say so, although Robert had been told by Jill a few hours before that she would not go because she did not have a costume to wear and had other plans for the night. Robert just completely forgot this when Anne asked him. When they arrive at the party, it turns out that Jill had changed her plans, and actually is at the costume party.

  Q: Was Robert permitted to tell Anne that Jill would be at the party?
  R1: Yes, because Jill did go to the party.
  R2: No, because Jill had told Robert that she would not go.
  R3: Other/Don't know.

Again, after each scenario, control questions were asked testing subjects' understanding of *truth* and *justification*.

4.2.2 Results

The distribution of responses in scenario B1 (N = 42) as well as B2 (N = 41) matches the distribution we obtained in Study 2B qualitatively. When Robert asserts a proposition that is false but which he is justified in believing, a substantial majority (81.0%) of the participants held that Robert is permitted to make the assertion. Only four participants stated that Robert is not permitted to do so, and a further four participants selected the 'Don't know' option. The results were reversed in scenario B2 in which Robert truly asserts that Jill is at the party despite not being justified. 58.5% of the participants chose 'No' when asked whether Robert was permitted to make that assertion, and 17.1% of the participants chose 'Yes'. Pearson's χ^2-test showed that the difference between the two scenarios was highly significant: $\chi^2 = 34.63$, $p < 0.001$, Cramer's V = 0.6. Focusing on those

participants who answered both control questions correctly, the distribution of the results was even more extreme (see Figure 3 for the percentages of the responses in both cases). For both scenarios, the data reveal significant results, way above chance level.

4.3 Discussion

Studies 2A and 2B show a remarkable consistency in the distribution of the responses. This close match allows us to discount an explanation of the results in terms of an artifact of the particular stories that we presented participants with. Instead, we are confident that the results we found in Study 2 are to be accounted for by the manipulations in regards to whether or not the protagonist's belief is true and justified. In scenarios A1 and B1, both Maria and Robert assert propositions that are not true, i.e. do not correspond with reality. However, both are justified in believing what they assert: Maria relied on a true memory that she had purchased a watch in the past. Robert was explicitly told by Jill that she would come to the party. The recorded results reveal that laypeople believe that a person who is justified in believing a certain proposition is also permitted to assert the proposition, even if the proposition is false – confirming the basic outcome of Study 1. Having gained similar response profiles in both kinds of scenarios, the effect is likely to be robust.

In scenarios A2 and B2, Maria and Robert assert propositions that are true but which they are not justified in believing, given their prior epistemic state. In both cases, the protagonists forgot information that would have led them to make different assertions. Their assertions turned out to be true, albeit through 'lucky' circumstances. The results from both experiments provide substantial evidence for (Justified Belief) and further question the plausibility of (True Belief). While the true belief account falsely predicts that most participants should approve of the assertions made because they were true, the justified belief account correctly predicts that the majority of the participants will reject that the protagonists were permitted to make the respective assertions. While the data from A2 and B2 are also consistent with the knowledge account of assertion, taking into account both studies, the only account that predicts the results across all studies is (Justified Belief).

5. GENERAL DISCUSSION

Studies 1 and 2 provide empirical evidence against both the knowledge account of assertion as well as the true belief account of assertion. Most empirical work on this subject
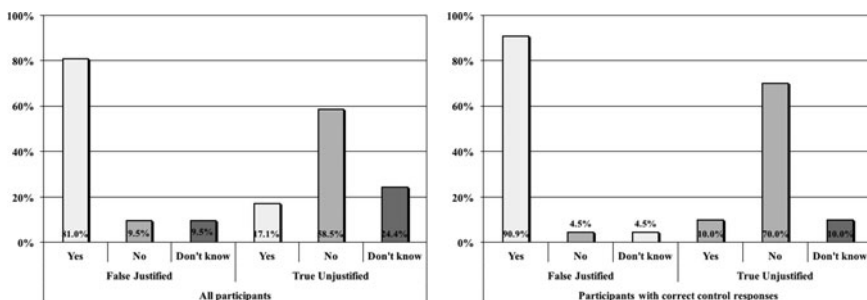

Fig. 3. Percentage of responses for Scenario B1 (False Justified) and Scenario B2 (True Unjustified).

matter has been conducted by Turri. As such, our own studies are inspired by his empirical approach to the problem – an approach we are in full agreement with. Turri (2016: 87) concludes that the sheer wealth of studies pointing towards the truth of the knowledge account settles the question on the norm of assertion. The results of our experiments tell a different story. While many experimental studies conducted by Turri are consistent with (Knowledge), they are also consistent with (Justified Belief). The key experiment that was supposed to tackle the issue of whether a justified belief may be sufficient for making an assertion, has been conducted in Turri (2013) and Turri (2017). Those experiments have yielded results that putatively rule out (Justified Belief). However, we have identified (in Section 2) two problems (awareness of error possibility and question type) with his experimental setup that questions the reliability of these results. In Section 3, we replicated Turri's original results but also demonstrated that once these issues were addressed, the results support more strongly (Justified Belief) and are inconsistent with (True Belief) and (Knowledge).[13]

Turri (2014) has drawn a distinction between good and permissible assertions that might be used to raise an objection against casting the question of the norm of assertion in terms of what one is permitted to say instead of what one should say. Put succinctly, one might worry that our studies have only shown that a justified false assertion is permissible but not that it is a good or proper assertion. While some authors have rejected an equivocal notion of the norms of assertion – see for example Lackey (2007) and Kvanvig (2011) – no consensus has so far emerged on this issue. However, the explanations we collected from the participants in Study 1 provide additional support that asking about when an assertion is *permitted* instead of when someone *should* make an assertion does not introduce any confounds but rather addresses the concerns about the *should* question that we discussed in Section 2. Importantly, the removal of either of the two problems we analyzed was enough to reverse the results, supporting (Justified Belief) in both cases. Thus, in order to show that truth is a necessary part of the norm of assertion and justified belief insufficient, one would need to show that both objections to the original set-up are misguided.

The results of Study 1 on their own, are, however, insufficient to show that justified belief is the norm of assertion. A disjunctive account seems equally apt to account for the results. More specifically, the data are consistent with the idea that a justified belief *or* a true belief are sufficient for a person to make an assertion. We therefore designed two sets of vignettes (Study 2) that pitted truth against justification. The outcome of Study 2A and 2B rule out the disjunctive account: the majority of the participants hold that a mere true belief does not permit a person to make an assertion if that person lacks justification for the belief. Thus, in summary we can state the following:

(True Belief) is not the norm of assertion, because *truth* is neither a necessary nor a sufficient condition for the norm of assertion.
(Knowledge) is not the norm of assertion, because *True Belief* is false.

---

13 There is an interesting question of whether non-binary accounts of justification norms as advanced by, for example, Goldberg (2015), McKinnon (2015) and Gerken (2017), can better account for the results of Study 1. Arguably, Maria is justified in believing that she has a Rolex watch, even if she is aware of an error possibility. A more nuanced view needs to consider the various degrees to which Maria is justified in her belief.

(True Belief or Justified Belief) is not the norm of assertion, because *justification* is necessary for being permitted to make an assertion.

(Justified Belief) is, ceteris paribus, the correct norm of assertion, because *justification* is both necessary and sufficient for being permitted to make an assertion.

Our experiments support (Justified Belief) as the norm of assertion. However, our experiments do not establish (Justified Belief) beyond a reasonable doubt. For once, we have not directly tested the importance of belief as a necessary component but primarily focused on the influence of justification. Second, there is one interesting hypothesis that we did not test for with our vignettes; it is the hypothesis that the norm of assertion is context-dependent. We know from many other empirical studies that people's intuitions are often context- and culture-dependent and not straightforwardly generalizable (see e.g. Machery *et al.*, 2013; Del Pinal and Reuter, 2016). More specifically, which norm of assertion holds in a given context might depend on (i) the authority/expertise of the agents as well as on (ii) the stakes involved.[14]

With respect to (i), we already argued in the introduction that our norms concerning when to trust the assertions of others should cohere with our norms of assertion. When we have to trust the assertions of others without any questions asked, we argued, (Knowledge) might be the appropriate norm of assertion. Here the asserter is solely responsible for ensuring that the agents listening learn something true and other agents are permitted to rely on the assertion without any questions asked. Indeed, when the asserter is an expert, e.g. a medical doctor, laypersons typically are in such a position. They have to trust the assertions and judgments of experts because they are not even in a position to assess possible answers to the given question or problem. Thus, they also cannot critically review the assertion of the expert (and typically the expert is aware of this). In such a context, it might very well be the case that the knowledge norm of assertion holds (respectively that the asserter's justification is required to be close to infallible).

In many contexts, however, we are in the position to critically review the assertions of our peers and their reasons for these assertions (at least to some extent) and the agents involved are typically aware of this. Thus, in such contexts it would make sense to decrease the requirements for being permitted to make an assertion (see Goldberg 2015, McKinnon 2015 and Gerken 2017, for more elaborate discussions of this point). At the same time the requirements for being permitted to trust the assertions of others could be increased. In such cases, agents can exchange and critically review each other's reasons. Consequently, we typically do not trust the assertions of our peers without any questions asked. We are also not surprised if we are asked to provide reasons for our claims. In these circumstances, (Justified Belief) might be the more appropriate norm. (This also seems to be the norm of assertion practiced at academic conferences; (Knowledge) seems to be more often violated than not at such events.)

With respect to (ii), it is possible that stakes influence what norms of assertion are appropriate. To be more exact, it is possible that stakes influence how strongly an agent needs to be justified to be permitted to make an assertion. Zollman (2015) has already argued to the same effect for norms concerning when to trust testimonial evidence.

---

14   Gerken (2017) discusses further factors that might have an impact on the norms of assertion, including urgency, available evidence, and alternative assertions.

He argued that whether or not an agent needs to be justified for being permitted to trust the testimony of others depends on our epistemic stakes and goals. According to Zollman, if the only epistemic goal is to minimize the number of false beliefs, then one should only trust those testifiers that one has strong reasons to trust, but if your goal is also to maximize the number of true beliefs, then it is better to be more trusting. The same can be expected for norms of assertion. If agents would only assert what they are absolutely certain of, they would further the satisfaction of the epistemic goal to minimize false beliefs. These agents would also possibly satisfy (Knowledge). However, if agents make assertions even if they are not absolutely certain but have good reasons to believe the claims, they further the satisfaction of the epistemic goal to maximize true beliefs at the risk of increasing the number of false beliefs. Such agents are more likely to violate (Knowledge). Of course, these considerations apply to epistemic stakes only. However, similar considerations might apply if the stakes are of a more pragmatic nature. Further studies are necessary to test the influence of the authority/expertise of the agent as well as the stakes involved.[15]

6. CONCLUSION

We started our investigation by highlighting the importance of knowing the norm of assertion. If knowledge were the norm of assertion, then people could straightforwardly believe honest assertions made by other parties. However, the knowledge account does not square with the empirical results that we obtained in two experiments presented in this paper. The central experiments favoring the knowledge account raise methodological concerns. Our first study has shown that once these concerns are addressed, collected responses from people reveal that truth and hence knowledge are not an integral part of the norm of assertion. In our second set of studies, we widened the scope of our investigation and provided evidence favoring the justified belief account of the norm of assertion. At least in many circumstances, people are permitted to make a certain claim whenever they are justified in believing that claim. No knowledge is required![16]

REFERENCES

Ben-Yami, H., Ben-Yami, M. and Ben-Yami, J. 2015. 'The False Belief Task Reconsidered.' *Budapest CEU Conference on Cognitive Development*, 160.
Del Pinal, G. and Reuter, K. 2016. 'Social Role Concepts, Commitments, and the Normative Dimension of Conceptual Representation.' *Cognitive Science*, 41: 477–501.
Douven, I. 2006. 'Assertion, Knowledge, and Rational Credibility.' *Philosophical Review*, 115: 449–85.
Gerken, M. 2013. 'Epistemic Focal Bias.' *Australasian Journal of Philosophy*, 91(1): 41–61.

---

15   Turri (2013) as well as Turri & Buckwalter (Forthcoming) tested the influence of stakes on people's responses in regards to the norms of assertion and did not find any significant effect.

16   We would like to thank the audiences in Basel, Berlin, Munich, Pittsburgh and Warsaw, at which previous versions of this paper have been presented. We are also grateful to Igor Douven, Anna-Maria Asunta Eder, Mikkel Gerken, Grzegorz Gaszczyk, Markus Kneer, Edouard Machery, Peter Pagin, Thomas Raleigh, Alex Wiegmann and an anonymous reviewer for *Episteme* for their very helpful comments and feedback.

—— 2017. *On Folk Epistemology: How We Think and Talk About Knowledge*. Oxford: Oxford University Press.

Gettier, E. 1963. 'Is Justified True Belief Knowledge?' *Analysis*, 23: 121–23.

Goldberg, S. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.

Kneer, M. Forthcoming. 'Norms of Assertion: Empirical Data.' *Cognition*.

Kvanvig, J. 2011. 'Norms of Assertion.' In J. Brown and H. Cappelen (eds), *Assertion: New Philosophical Essays*, pp. 233–50. Oxford: Oxford University Press.

Lackey, J. 2007. 'Norms of Assertion.' *Noûs*, 41: 594–626.

Machery, E., Mallon, R., Nichols, S. and Stich, S. 2013. 'If Intuitions Vary, Then What?' *Philosophy and Phenomenological Research*, 86: 618–35.

Maitra, I. 2011. 'Assertion, Norms, and Games.' In J. Brown and H. Cappelen (eds), *Assertion: New Philosophical Essays*, pp. 277–96. Oxford: Oxford University Press.

Marsili, N. and Wiegmann, A. Ms. 'Should I Say That? An Experimental Investigation into the Norm of Assertion.'

McKinnon, R. 2015. *The Norms of Assertion: Truth, Lies, and Warrant*. London: Palgrave Macmillan.

Nagel, J., San Juan, V. and Mar, R. 2013. 'Lay Denial of Knowledge for Justified True Beliefs.' *Cognition*, 129: 652–61.

Pagin, P. 2015. 'Assertion.' In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/archives/win2016/entries/assertion/.

—— 2016. 'Problems with Norms of Assertion.' *Philosophy and Phenomenological Research*, 93(1): 178–207.

Perner, J., Leekam, S. R. and Wimmer, H. 1987. 'Three Year Olds' Difficulty With False Belief: The Case for a Conceptual Deficit.' *British Journal of Developmental Psychology*, 5(2): 125–37.

Turri, J. 2013. 'The Test of Truth: An Experimental Investigation of the Norm of Assertion.' *Cognition*, 129: 279–91.

—— 2014. 'Knowledge and Suberogatory Assertion.' *Philosophical Studies*, 167: 557–67.

—— 2015. 'Knowledge and the Norm of Assertion: A Simple Test.' *Synthese*, 192(2): 385–92.

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

—— 2017. 'The Distinctive "Should" of Assertability.' *Philosophical Psychology*, 30(4): 477–85.

Turri, J. and Buckwalter, W. Forthcoming. 'Descartes's Schism, Locke's Reunion: Completing the Pragmatic Turn in Epistemology.' *American Philosophical Quarterly*.

Williamson, T. 1996. 'Knowing and Asserting.' *Philosophical Review*, 105: 489–523.

—— 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Zollman, K. 2015. 'Modeling the Social Consequences of Testimonial Norms.' *Philosophical Studies*, 172(9): 2371–83.

KEVIN REUTER is a lecturer and post-doctoral researcher at the Institute of Philosophy, University of Bern. His research interests are in philosophy of mind and psychology, cognitive and social sciences, and experimental philosophy. He is also the principal investigator of the research project "The Conceptual Space of the Affective Mind" funded by the Swiss National Science Foundation.

PETER BRÖSSEL holds the post of Junior-Professor for the Philosophy of Perception and Knowledge and director of the Emmy Noether research group "From Perception to Belief and Back Again" at the Department of Philosophy II and the Center for Mind, Brain, and Cognitive Evolution at the Ruhr-University Bochum. His research interests focus on epistemology, philosophy of science, and philosophy of mind and language.