

Knowledge and Assertion in “Gettier” Cases*

John Turri

Philosophy Department and Cognitive Science Program
University of Waterloo
200 University Avenue West
Waterloo, Ontario N2L3G1, Canada
E: john.turri@gmail.com

Abstract: Assertion is fundamental to our lives as social and cognitive beings. By asserting we share knowledge, coordinate behavior, and advance collective inquiry. Accordingly, assertion is of considerable interest to cognitive scientists, social scientists, and philosophers. This paper advances our understanding of the norm of assertion. Prior evidence suggests that knowledge is the norm of assertion, a view known as “the knowledge account.” In its strongest form, the knowledge account says that knowledge is both necessary and sufficient for assertability: you should make an assertion if and only if you know that it is true. The knowledge account has been rejected on the grounds that it conflicts with our ordinary practice of evaluating assertions. This paper reports four experiments that address an important objection of this sort, which focuses on a class of examples known as “Gettier cases.” The results undermine the objection and, in the process, provide further evidence for the knowledge account. The findings also teach some important general lessons about intuitional methodology and the curation of genres of thought experiment.

Keywords: norms; knowledge; assertion; social cognition; philosophical method

* This is the penultimate version of a paper to appear in *Philosophical Psychology*. Please cite the final, published version if possible.

Introduction

Assertion is ubiquitous, unavoidable and extremely important in human society. Our individual and collective well being often depends on it: it is the central means of sharing information and coordinating actions. Unsurprisingly, then, researchers have taken a keen interest in the norm of assertion: under what conditions should you make an assertion? In recent decades, there has been a considerable amount of philosophical work investigating the norm of assertion (Unger, 1975; Bach & Harnish, 1979; Williamson, 1996, 2000; Reynolds, 2002; Weiner, 2005; Douven, 2006; Lackey, 2007; Hill & Schechter, 2007; Schaffer, 2008; Turri, 2010a; Turri, 2010b; Turri, 2014a), and more recently cognitive scientists have begun examining the issue too (e.g. Turri, 2013a; Turri, 2015a; Turri, in press).

What is the norm of assertion? According to *the knowledge account*, knowledge is the norm of assertion. More specifically, knowledge is the norm of the actual social practice of assertion. In its strongest form, the knowledge account says that knowledge is both necessary and sufficient for assertability: you should make an assertion if and only if you know that it is true. A wealth of theoretical and empirical evidence supports the knowledge account (see Turri, 2016 for an overview). However, critics charge that the knowledge account mischaracterizes our ordinary practice of evaluating assertions. To support this charge, critics have tried to produce counterexamples to the knowledge account. One common strategy is to try to identify cases where, intuitively, assertability is present even though knowledge is lacking.

Researchers investigating the norm of assertion agree that the project is, at least in large part, empirical. As critics themselves note, “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.” This implies that “any proposal made in the course of” the investigation is “subject to the exact same standards of evaluation as are employed in the empirical sciences generally.” In short, hypotheses in this area “must face the linguistic data” (Douven, 2006, p. 450). This is a fair point because if knowledge is the norm of assertion, then full competence in the practice of assertion requires sensitivity to the connection between knowledge and assertability. This does not imply that skilled practitioners have an explicit, articulable theory about what the norm is, or that they will answer “knowledge” if asked “what is the norm of assertion?” The assumption is not that skilled practitioners tend to be good theorists; rather, it’s that patterns in their concrete, first-order judgments manifest their skill in applying the relevant rules (compare Chomsky, 1957; Noveck & Sperber, 2004).

Sometimes critics focus on cases of reasonable false assertion (Douven, 2006; Lackey, 2007; Hill & Schechter, 2007; Kvanvig, 2009). The idea here is that it is sometimes reasonable to believe false propositions, and if it is reasonable for you to believe a false proposition, then intuitively there is no sense in which you should not assert it. But knowledge requires truth, so you can’t know any such proposition and, hence, knowledge isn’t the norm of assertion. However, this argument fails because the critics have mischaracterized how reasonable false assertions are actually evaluated (Turri, 2013a; Turri & Blouw, 2015). More specifically, in experiments where hundreds of ordinary adult speakers were asked to evaluate whether someone should make an assertion, the assertion’s truth-value had a powerful influence on whether people judged that the agent should make the assertion. When the assertion would be true, virtually everyone said that the agent should make the assertion. But when the assertion would be false, the vast majority said that she should not make the assertion. This same basic pattern persisted

when people were questioned in different ways and across other differences that can influence evaluative judgments and social cognition.

Other times critics focus on cases of “selfless assertion” (Lackey, 2007, pp. 598 ff.; Pritchard, 2014, p. 160; Wright, 2014, p. 255). A selfless assertion is an assertion that allegedly has two crucial features. First, it’s an assertion that, intuitively, the agent should make. Second, we naturally interpret the agent as neither believing nor, as a result, knowing the proposition asserted. In other words, the idea here is that sometimes people should assert propositions that they don’t believe. But knowledge requires belief, so they don’t know these propositions and, hence, knowledge isn’t the norm of assertion. However, this argument fails too because in fact “selfless assertors” are readily interpreted as both believing and knowing the propositions they assert (Turri 2015b). More specifically, in experiments where hundreds of ordinary adult speakers were asked to evaluate cases of selfless assertion from the literature, the vast majority judged that the agent should assert, believes, and knows the proposition in question. Thus, such cases cannot play the destructive role that critics have assigned them.

For both reasonable false assertions and selfless assertions, then, critics have misunderstood or mischaracterized the ordinary, intuitive evaluation of the cases. But one final popular example among critics has not been tested: “Gettiered assertions.” The idea here is that it is sometimes reasonable to believe true propositions that you nevertheless fail to know, due to objectionable forms of luck. These are often called “Gettiered beliefs,” named after Edmund Gettier, the philosopher who sparked discussion of such examples in the mid-twentieth century (Gettier, 1963; for some pre-1963 history of such cases, see Matilal, 1986, pp. 135-7; Chisholm, 1989: pp. 92-3). According to critics, Gettiered beliefs obviously fall short of knowledge but, intuitively,

there is no sense in which you should not assert them. Hence, critics argue, knowledge isn't the norm of assertion (e.g. Hill & Schechter, 2007; Lackey, 2007; Brown, 2008; Smithies, 2012; Smith, 2012; Coffman, 2014). However, by this point the critics' poor empirical track record should give us pause.

“Gettier cases” come in many varieties. Here I'll focus on two basic types frequently mentioned in the assertion literature. There might be no theoretically neutral way of describing the structure of these cases, but I'll try to remain as theoretically neutral as possible. But before proceeding it's worth clarifying something important. Although many philosophers have said — over dissent from other philosophers (e.g. Sartwell, 1992; Lycan, 2006; Turri, 2012; Hetherington, 2013) and in opposition to some central tendencies in commonsense epistemology (e.g. Starmans & Friedman 2012; Colaco, Buckwalter, Stich & Machery, 2014) — that “Gettiered beliefs” of all sorts are obviously not cases of knowledge, I do not simply assume that a “Gettier case” essentially involves ignorance. Instead, I am primarily interested in patterns of judgments about knowledge and assertability in cases with certain structures that have been put forward as compelling counterexamples to the knowledge account. These happen to be labelled “Gettier cases,” so I adopt that term for convenience. Critics might be wrong about whether these are viewed as cases of ignorance, or about whether they are naturally viewed as cases of assertability, about both matters, or about neither.

On the one hand, there are “environmental threat” or “fake barn” cases (Goldman, 1976, pp. 772-3; Goldman, 2009, p. 79, n. 5, crediting Carl Ginet). This is the most popular type of case among critics of the knowledge account. In an environmental threat case, the agent believes that something is true because she directly perceives it. If that were the end of the story, then

intuitively she would know that the proposition is true. But it turns out that the agent is in an environment where her perceptual evidence could very easily have been misleading and led her to form a false belief. Intuitively, many philosophers claim, this real and very near possibility of error prevents the agent from knowing (e.g. Goldman, 1976; Sosa, 1991, pp. 238-9; Neta & Rohrbaugh, 2004, p. 401; Pritchard, 2005, pp. 161-2; Kvanvig, 2008, p. 274). For example, suppose Sarah looks out her car window and sees a roadside barn as she drives along. Everything about Sarah and the barn is normal. But Sarah doesn't realize that the area she is driving through is being used as a movie set and the set designers have constructed many fake-barn façades that look just like real barns. Sarah is looking at the one real barn among all the nearby fakes. Clearly Sarah doesn't know that it's a barn, the critic claims, but surely Sarah should, if asked, say that it's a barn.

On the other hand, there are “explanatory disconnect” or “apparent evidence” cases (the latter label is due to Starmans & Friedman, 2012). In an apparent evidence case, an agent believes a true proposition based on good but fallible evidence. If that were the end of the story, then presumably he would know that the proposition is true. But it turns out that the agent's evidence is misleading and his belief is made true by something completely unrelated to his evidence. Intuitively, the unexpected explanatory disconnect between evidence and truth prevents the agent from knowing. For example, suppose that Angelo is in the forest during deer hunting season. Two very loud, sharp bangs ring out nearby. Angelo judges that somebody is hunting deer nearby. And there is somebody hunting deer nearby. But the bangs Angelo heard were just backfire from a vehicle, and his belief is true because a camouflaged hunter is stalking a deer nearby with bow-and-arrow, silent and unseen. Clearly Angelo doesn't know that

someone is hunting nearby, the critic claims, but surely Angelo should, if asked, say that someone is hunting nearby.

The experiments reported below test both types of case: explanatory disconnect and environmental threat. In each experiment, I collect judgments about two critical issues: whether the agent knows and whether the agent should assert. If critics are correct, then the central tendency will be for people to deny knowledge but attribute assertability. However, the results show that the central tendency is for judgments about knowledge and assertability to remain united. The findings also suggest some important general lessons for philosophers about intuitional methodology and the curation of genres of thought experiment.

Experiment 1

This experiment provides a simple initial test of assertability judgments in two different types of Gettier case: an environmental threat case and an explanatory disconnect case.

Method

Participants. Two hundred U.S. residents were tested (aged 18-72 years, mean age = 33 years; 94% reporting English as a native language; 76 female). Participants were recruited and tested online using Amazon Mechanical Turk and Qualtrics and compensated \$0.35 for approximately 2 minutes of their time. Repeat participation was prevented within and across experiments.

Materials and Procedure. Participants were randomly assigned to one of four conditions in a 2 (Luck: Control/Experimental) \times 2 (Structure: Threat/Disconnect) between-subjects design. Participants read a single story. The story for the Experimental Threat condition featured a true belief based on uninterrupted direct perceptual access to the truth-maker; it is a version of the fake-barn example from the Introduction. The story for Experimental Disconnect condition

featured a true belief based on good, fallible and misleading evidence (“apparent” evidence); it is a version of the deer-hunting example from the Introduction. The Control Threat and Control Disconnect conditions featured closely matched stories that are similar in length and complexity to their experimental counterparts, but they lack the “bad luck” that is widely thought to undermine knowledge in Gettier cases.

Following prior methodological suggestions for testing intuitions about Gettier cases, the stories were presented in three stages to participants. After reading each of the first two parts of the story, participants were asked whether the agent saw or heard the relevant item (response options rotated randomly). These questions were asked to prime participants to attend to these important details of the story; there is some evidence that this priming is effective (Turri, 2013b, section 9). The Control and Experimental conditions differed only in some crucial details presented in the story’s second stage. Here is how the story proceeded for the Threat conditions:

(Stage 1) Sarah is driving with her son down the highway. Sarah looks out the window of her car and sees a red barn near the road. [Question: Sarah sees a barn. (Yes/No)]

(Stage 2: Control/Experimental) Sarah doesn’t realize that the countryside she is driving through is currently being used as the set of a film, and that the set designers have constructed many [cheap barns/fake barn facades] in this area that look as though they are [expensive/real] barns. ¶¹ Despite all the [cheap barns/fakes] around, Sarah is in fact looking at the one [expensive/real] barn in the area. [Question: Sarah sees a barn. (Yes/No)]

1 Indicates a paragraph break on the participant’s screen.

(Stage 3) Sarah's son points to the barn and says, "Mom, I have to do a report on barns for my social studies class. Is that a barn?"

In order to measure judgments about knowledge and assertability, participants were then instructed, "Please select the option that best describes Sarah." They had four options (rotated randomly; participants did not see the numerical labels):

1. She knows that it's a barn, and she should tell her son that it's a barn.
2. She knows that it's a barn, and she should not tell her son that it's a barn.
3. She does not know that it's a barn, and she should tell her son that it's a barn.
4. She does not know that it's a barn, and she should not tell her son that it's a barn.

Here is how the story proceeded for the Disconnect conditions:

(Part 1) Angelo is camping with his daughter in a wooden cabin at the edge of the forest. As they settle in to sleep for the night, the daughter has her headphones on and Angelo is reading near the window. Angelo hears two very loud, sharp bangs ring out in the forest behind the cabin. It is deer-hunting season. [Question: Angelo heard gunshots. (Yes/No)]

(Part 2: Control/Experimental) Angelo doesn't realize that the [gunshots/bangs] he heard were [special/not gunshots]. They [came from a daredevil who is hunting deer while driving an all-terrain vehicle/were just backfire from an all-terrain vehicle being driven by a daredevil] along the forest trail. ¶ Angelo doesn't realize something else too. An honorable veteran deer hunter is silently stalking with a bow and arrow nearby. He has a deer in his sights. The hunter is camouflaged and

no one, not even the deer, has seen or heard him. [Question: Angelo heard gunshots. (Yes/No)]

(Part 3) Angelo's daughter takes off her headphones and asks, "Dad, what's going on? Is somebody hunting deer nearby?"

Participants were then instructed, "Please select the option that best describes Angelo." They had four options.

1. He knows that someone is hunting nearby, and he should tell his daughter that someone is hunting nearby.
2. He knows that someone is hunting nearby, and he should not tell his daughter that someone is hunting nearby.
3. He does not know that someone is hunting nearby, and he should tell his daughter that someone is hunting nearby.
4. He does not know that someone is hunting nearby, and he should not tell his daughter that someone is hunting nearby.

Participants filled out a brief demographic questionnaire after testing.

Results

The priming questions effectively guided participants to attend to the relevant details, with participants answering correctly 93% of the time.² A preliminary multinomial logistic regression revealed main and interaction effects on response to the test question, so I will analyze the results for each type of case separately.

² These were not intended as comprehension questions. Still, it is worth noting that eliminating participants who gave "incorrect" answers yields the same basic pattern of results reported below.

For the Threat cases, there was no effect of Luck on response to the test question, $\chi^2(3, 100) = 4.14, p = .246, n.s.$ An overwhelming majority answered that Sarah knows the proposition and should assert it: Control Threat, 94%, binomial test, $p < .001$, all tests two-tailed, test proportion for all binomial tests reported for this experiment = .25; Experimental Threat, 83%, binomial test, $p < .001$. Only a very small number of participants answered that Sarah does not know the proposition and should assert it: Control Threat, 4%, binomial test, $p < .001$; Experimental Threat, 8%, binomial test, $p = .002$. (See Fig. 1.)

For the Disconnect cases, there was an effect of Luck on response to the test question, $\chi^2(3, 100) = 21.78, p < .001$. The magnitude of the difference in frequencies was large, Cramer's $V = .467$. In Control Disconnect, a majority of participants (63%) answered that Angelo knows the proposition and should assert it, binomial test, $p < .001$. In Experimental Disconnect, a majority (57%) answered that Angelo does not know the proposition and should not assert it, binomial test, $p < .001$. Only a small number of participants answered that Angelo does not know the proposition and should assert it: Control Disconnect, 14%, binomial test, $p = .053$; Experimental Disconnect, 20%, binomial test, $p = .238$.

It is also important to consider how frequently participants link knowledge and assertability. Whether participants said that the agent knows and should assert, on the one hand, or does not know and should not assert, on the other, the response is fully consistent with the knowledge account of assertion. Accordingly, I analyzed how frequently participants chose one or the other of these two options that keep knowledge and assertability together. For simplicity, call either sort of answer a *unified response*. A binary logistic regression revealed no main or interaction effects of Structure or Luck on whether participants offered a unified response; the

model was statistically insignificant, $\chi^2(3, 200) = 7.18, p > .05$. Overall, 85% (170 of 200) of participants offered a unified response, $\chi^2(1, 200) = 98.0$.




Fig. 1. Experiment 1. Number of participants in the four conditions who selected the various options as best describing the agent in the story. “K&S” = the agent knows the proposition and should say it. “K&~S” = the agent knows the proposition and should not say it. “~K&S” = the agent does not know the proposition and should say it. “~K&~S” = the agent does not know the proposition should not say it.

Discussion

The results are inconsistent with many researchers’ claims about cases like the ones tested. They claim that the Experimental cases tested above are intuitive counterexamples to the knowledge account. In particular, they claim that the Experimental cases are intuitive examples where someone should assert a proposition that she or he doesn’t know. But the responses observed here tell a very different story. In the Experimental Threat condition, participants

overwhelmingly judged that the agent both knew the proposition and should assert it. In the Experimental Disconnect condition, a strong majority judged that the agent neither knew the proposition nor should assert it. These results are fully consistent with the knowledge account and show that, in principle, such cases pose no problem for the account.

Experiment 2

In Experiment 1, I tested two different types of case often labeled “Gettier cases.” But there are other possible variations on the Gettier theme. For example, in the story for the Experimental Disconnect condition, the natural interpretation of the case is that Angelo is right that someone is hunting deer nearby, and that he came to believe this because he heard the loud bangs and interpreted them as gunshots. (Recall that Angelo is out in the forest during deer hunting season.) But the loud bangs are not gunshots; they’re backfire from a nearby vehicle instead. So Angelo is fundamentally wrong about the nature of his evidence. However, it could be argued that this is not the hardest test of the knowledge account. Angelo doesn’t need to be fundamentally wrong about the nature of his evidence in order for there to be an unexpected explanatory disconnect between his evidence and what makes his belief true. For example, it could be that Angelo does hear gunshots but they are from someone taking target practice rather than someone actually hunting. Perhaps a case structured this way will pose a greater challenge to the knowledge account of assertion. This experiment investigates this possibility.

Method

Participants. Two hundred twenty-five new participants were tested (aged 18-70 years, mean age = 32 years; 93% reporting English as a native language; 85 female). Participants were recruited and tested the same way as in Experiment 1.

Materials and Procedure. Participants were randomly assigned to one of three conditions — Control, Disconnect Light (“Light” for short), and Disconnect Heavy (“Heavy” for short) — in a between-subjects design. Participants all read a single story very similar to the story used in the Disconnect conditions in Experiment 1. The stories for the three conditions differed in some crucial details in the first paragraph of the second stage. Otherwise, the stories, questions, and procedures were exactly the same as in the Disconnect conditions in Experiment 1. Here is the distinguishing paragraph for each condition.

(Control) Angelo doesn’t realize that the gunshots he heard were fired by a very famous hunter. The shooter is an off-duty marine hunting deer in the forest. He enjoys hunting deer.

(Light) Angelo doesn’t realize that the gunshots he heard were not fired by a hunter at all. The shooter is an off-duty marine taking target practice in the forest. He never hunts deer.

(Heavy) Angelo doesn’t realize that the bangs he heard were not gunshots at all. They are just backfire from a vehicle being driven by an off-duty marine in the forest. He never hunts deer.

Results

The priming questions again effectively guided participants to attend to the relevant details, with participants answering correctly 90% of the time.³ There was an effect of condition on response to the test question, $\chi^2(6, 225) = 33.91, p < .001$, Cramer’s $V = .275$. (See Fig. 2.) Pairwise

3 Again, eliminating participants who gave “incorrect” answers yields the same basic pattern of results reported below.

comparisons revealed that response to the test question differed significantly between Control and Light, $\chi^2(3, 150) = 12.85, p = .005$, Cramer's $V = .293$; it differed significantly between Control and Heavy, $\chi^2(3, 149) = 24.34, p < .001$. Cramer's $V = .404$; and it differed significantly between Light and Heavy, $\chi^2(3, 151) = 11.27, p = .01$, Cramer's $V = .273$.

In Control, a majority (64%) answered that Angelo knows the proposition and should assert it, binomial test, $p < .001$, test proportion for all binomial tests reported in this experiment = .25. Only a very small number of participants (9%) answered that Angelo does not know the proposition and should assert it, binomial test, $p = .001$.

In Light, a plurality (39%) answered that Angelo knows the proposition and should also assert it, binomial test, $p < .001$. A nontrivial number of participants (30%) answered that Angelo does not know the proposition and should assert it, but this did not differ from what could be expected by chance, binomial test, $p = .352$.

In Heavy, a majority (51%) answered that Angelo does not know the proposition and should not assert it, binomial test, $p < .001$. A small number of participants (19%) answered that Angelo does not know the proposition and should assert it, which did not differ from what could be expected by chance, binomial test, $p = .254$.

Again I analyzed how frequently participants offered a unified response for knowledge and assertability (i.e. answered that the agent knows and should assert, on the one hand, or does not know and should not assert, on the other). There was an effect of condition on how frequently participants offered a unified response, $\chi^2(2, 225) = 4.72, p = .023$. The magnitude of the difference in frequencies was small, Cramer's $V = .183$. In Control, 85% of participants offered a unified response, $\chi^2(1, 74) = 36.54, p < .001$; in Light, 66% did, $\chi^2(1, 76) = 7.58, p = .006$; in

Heavy, 75% did, $\chi^2(1, 75) = 18.25, p < .001$. Overall, 75% of participants offered a unified response, $\chi^2(1, 225) = 56.75, p < .001$.




Fig. 2. Experiment 2. Number of participants selecting the options as best describing the agent in the story. “K&S” = the agent knows the proposition and should say it. “K&~S” = the agent knows the proposition and should not say it. “~K&S” = the agent does not know the proposition and should say it. “~K&~S” = the agent does not know the proposition should not say it.

Discussion

The results replicate the main findings from Experiment 1. Responses were once again fully consistent with the knowledge account. The current results also generalized the main finding to a Gettier case with an importantly different structure.

The current results sharpen a lesson from Starmans and Friedman’s important work on knowledge attributions (Starmans & Friedman, 2012). Starmans and Friedman distinguished

between Gettier cases involving “apparent evidence” and “authentic evidence.” Authentic evidence is, roughly, evidence that makes a belief true when based on it. Apparent evidence, by contrast, merely appears to make a belief true when based on it. Starmans and Friedman discovered that people are much less likely to attribute knowledge in apparent evidence cases. In this experiment, people in the Light Disconnect and Heavy Disconnect conditions attributed knowledge at different rates. But it seems that neither condition featured an authentic evidence case. Assuming that the apparent/authentic distinction is exhaustive, then, the present result demonstrate a relevant difference entirely within the class of apparent evidence cases. In particular, the results suggest that the more fundamentally an agent is misled about the nature of his evidence, the less likely people are to attribute knowledge.

Experiment 3

To further test the robustness of the findings above, I tested people’s judgments about knowledge and assertability with a completely different “apparent evidence” cover story and completely different dependent measures. Researchers have suggested that the contrast “knows/doesn’t know” is less meaningful to people than the contrast “know/only thinks” or “knows/only thinks he knows” (Starmans & Friedman, 2012; Turri, 2013b; see also Weinberg, Nichols & Stich, 2001; Cullen, 2010). In particular, it seems that a contrast with “only thinks” allows people to acknowledge how things seem from the agent’s perspective while, at the same time, indicate that the agent’s perspective is incorrect. It’s also possible, especially when considering Gettier cases, that the agent’s state is *much more like* knowing than ignorance, in which case people might answer that the agent “knows” because it’s the best option available to them, even though they don’t quite think that the agent knows. In light of this, I used a “thinks” contrast for both the

knowledge and assertability judgments in this experiment.

Method

Participants. One hundred fifty-one new participants were tested (aged 18-65 years, mean age = 33 years; 94% reporting English as a native language; 55 female). Participants were recruited and tested the same way as in earlier experiments.

Materials and Procedure. Participants were randomly assigned to one of two conditions, Control and Experimental, in a between-subjects design. Participants read a single story about Geno. Geno's mother is completing a home improvement project and she needs a set of metric wrenches. Her old set is lost so Geno goes to the hardware store to buy a new set. In the story for each condition, Geno buys a set of wrenches and puts them in the garage. In the Control story, it turns out that Geno bought carbon steel metric wrenches rather than, as his mother wanted, stainless steel. In the Experimental story, it turns out that Geno actually bought Imperial wrenches rather than, as his mother wanted, metric. In each story, it also turns out that the lost set of metric wrenches are also in the garage, although they will not be found. Geno tells his mother that there are metric wrenches in the garage, at which point she discovers that the wrenches he bought are either the wrong sort of steel (carbon in Control), or in the wrong units (Imperial in Experimental).

As in previous experiments, the story was presented in three stages. A priming question was asked at the end of the first two stages. A knowledge question was also asked at the end of the second stage. Then a retrospective assertability question was asked at the end of the third stage. Here is how the story and questioning proceeded.

(Stage 1) Geno's mother is planning a home improvement project. She needs a set

of metric wrenches (measured in millimeters), but she lost her old set and no one can find them now. So Geno buys a set of wrenches from the display in the hardware store. Then he puts them in the garage. [Question: Geno put metric wrenches in the garage. (Yes/No)]

(Stage 2: Control/Experimental) Geno didn't notice that the display at the hardware store was for [*carbon* steel metric wrenches (made with carbon)/stainless steel *Imperial* wrenches (measured in inches)]. But his mother asked for [*stainless* steel metric wrenches/*metric* wrenches]. ¶ Geno didn't notice something else. His mother's old set of metric wrenches is in the garage. They are under some old scrap metal in a garbage can where they will never be found. [Question: Geno put metric wrenches in the garage. (Yes/No)]

1. Geno _____ that there are metric wrenches in the garage. (knows/only thinks he knows)

(Stage 3: Control/Experimental) Geno tells his mother that there are metric wrenches in the garage. She finds the wrenches he bought and says, "I said to get stainless steel metric wrenches. These are [*carbon* steel metric wrenches/stainless steel *Imperial* wrenches]."

2. Geno _____ told his mother that there are metric wrenches in the garage. (should have/only thought he should have)

Results

The priming questions again effectively guided participants to attend to the relevant details, with

participants answering correctly 88% of the time.⁴ In order to measure judgments about knowledge and assertability, I analyzed the results for the four different combinations of answers that participants could select (the participant's *combined response*): (1) Geno knows and should have asserted, (2) Geno knows and only thought he should have asserted, (3) Geno only thinks he knows and should have asserted, (4) Geno only thinks he knows and only thought he should have asserted. These correspond to the same four possibilities that, in earlier experiments, participants chose among as explicit combinations. But this time the knowledge and assertability judgments occurred separately on different screens.

There was an effect of condition on combined response, $\chi^2(3, 151) = 24.16, p < .001$. (See Fig. 3.) The magnitude of the difference in frequencies was medium, Cramer's $V = .400$. In the Control condition, a majority of participants (53%) answered that Geno knows the proposition and should have asserted it, binomial test, $p < .001$, test proportion for all binomial tests reported in this experiment = .25. A nontrivial number of participants (25%) answered that Geno only thinks he knows the proposition and should have asserted it, but this did not differ from what could be expected by chance, binomial test, $p = .99$. In the Experimental condition, a plurality of participants (44%) answered that Geno only thinks he knows the proposition and only thought that he should have asserted it, binomial, $p < .001$. A nontrivial number of participants (23%) answered that Geno only thinks he knows the proposition and should have asserted it, but again this did not differ from what could be expected by chance, binomial test, $p = .744$. There was no effect of condition on whether participants chose this combination of assertability without

4 Again, eliminating participants who gave "incorrect" answers yields the same basic pattern of results reported below.

knowledge, $\chi^2(1, 151) = 0.11, p = .849$.

Again I analyzed how frequently participants offered a unified response for knowledge and assertability (i.e. ascribed both knowledge and assertability, or denied both). There was no effect of condition on whether participants offered a unified response, 65% in each condition, $\chi^2(1, 151) = 0.01, p = 1$. Overall, 65% of participants offered a unified response, $\chi^2(1, 151) = 13.41, p < .001$.




Fig. 3. Experiment 3. Number of participants in the two conditions who selected the various combinations as best describing the agent in the story. “K&S” = the agent knows the proposition and should say it. “K&~S” = the agent knows the proposition and should not say it. “~K&S” = the agent does not know the proposition and should say it. “~K&~S” = the agent does not know the proposition should not say it.

Discussion

The results replicate the main findings from earlier experiments. Once again responses were consistent with the knowledge account. The central tendency in the control case was to attribute both knowledge and assertability, while the central tendency in the Gettier case was to deny both knowledge and assertability. This same pattern emerged while using a completely new cover story and a completely different dependent measure. Whereas in earlier experiments participants registered knowledge judgments and prospective assertability judgments in tandem, here participants registered knowledge judgments and retrospective assertability judgments

separately. This time the choices for the knowledge judgments and assertability judgments were different too, incorporating “thinks” contrasts to guard against potential projective or overly charitable readings. Despite these considerable differences in content and procedure, overall the same basic pattern of responses emerged.

Experiment 4

To further test the robustness of the findings above, I did two things in this experiment. First, I tested people’s judgments about knowledge and assertability with another new “apparent evidence” story based on a famous thought experiment in the literature. In this story, an office worker is misled about whether a co-worker owns a certain type of car (Lehrer, 1965, p. 169ff). The original story in the literature features intentional deception. However, intentional deception is inessential to the basic structure of an “apparent evidence” Gettier case, and the presence of deception affects social judgments and reasoning independently of any Gettier factors (Schul, Mayo, Burnstein & Yahalom, 2007). So I treated deception as a factor in the experimental design to see whether it affects intuitions in Gettier cases.

Second, I tested people’s judgments using different procedures. Until now I’ve followed prior research by presenting cases in three stages and including priming questions to encourage people to keep track of relevant details. I have also used different contrasts for various probes, including “only thinks” contrasts, which prior work has also suggested is more meaningful to people when making knowledge judgments. But some might wonder whether this approach is too didactic. In particular, some might wonder whether it unfairly nudges people to give answers favorable to the knowledge account of assertion. So I treated procedure as a factor in the experimental design to test this worry.

Method

Participants. Four hundred fifty-five new participants were tested (aged 18-77 years, mean age = 31 years; 95% reporting English as a native language; 165 female). Participants were recruited and tested the same way as in earlier experiments.

Materials and Procedure. Participants were randomly assigned to one of six conditions in a 3 (Story: Control/Non-deceptive/Deceptive) × 2 (Procedure: Guided/Unguided) between-subjects design. The story featured Douglas and Kelly. Kelly does and says things that lead Douglas to think that she owns a Ferrari. In the Control stories, everything is as it seems and Kelly does own a Ferrari. In the Non-deceptive stories, Kelly does not own a Ferrari but no intentional deception is involved. In the Deceptive stories, Kelly does not own a Ferrari and intentional deception is involved. In Guided procedural conditions, participants encountered the material and registered knowledge and assertability judgments just as in Experiment 3: they read the story in stages, answered priming questions (“Kelly owns a Ferrari” [Yes/No]), then answered a dichotomous knowledge question contrasting “knows/only thinks he knows,” then answered a dichotomous assertability question contrasting “should/only thinks he should.” In Unguided procedural conditions, participants read the story all at once and then, on the same screen, answered a knowledge question contrasting “knows/doesn’t know” contrasts and an assertability question contrasting “should/should not.” The knowledge and assertability probes were rotated randomly so that participants had an equal chance of receiving either one first. In all conditions, response options were rotated randomly.

Here is the text of the stories, which, as before, differ in crucial details of the second paragraph:

Douglas observes his co-worker, Kelly, arrive at work driving a new Ferrari.

She is clearly very happy with the car. A little later, while Kelly is on the phone with her cousin, Douglas hears her say, “I really love my new car! It might have cost a lot, but it is completely worth it.”

(Control) Douglas doesn't realize that the Ferrari is only one of Kelly's cars. Her other car is with her cousin. Kelly's other car is actually a Toyota hybrid. Just for fun, the cousin borrowed the Toyota for the day. But Kelly also had another reason for lending it: she wants to convince her cousin to get a hybrid. Kelly has always cared about the environment, and she thinks it will make a difference. And her plan worked: her cousin was completely convinced. It was a good thing.

(Non-deceptive) Douglas doesn't realize that the Ferrari is not Kelly's car. It belongs to Kelly's cousin. Kelly's new car is actually a Toyota hybrid. Just for fun, the cousins switched cars for the day. But Kelly also had another reason for switching: she wants to convince her cousin to get a hybrid. Kelly has always cared about the environment, and she thinks it will make a difference. And her plan worked: her cousin was completely convinced. It was a good thing.

(Deceptive) Douglas doesn't realize that the Ferrari is not Kelly's car. It belongs to Kelly's cousin. Kelly's new car is actually a Toyota hybrid. Just for fun, the cousins switched cars for the day. But Kelly also had another reason for switching: she wants to trick Douglas into thinking that she owns a Ferrari. Kelly has always had a crush on Douglas, and she thinks it will impress him. And her plan worked: Douglas was completely tricked. It was a total lie.

Douglas doesn't realize something else. His great grand-uncle recently died.

In his will, the uncle left Douglas a Ferrari. So Douglas now owns a Ferrari. But Douglas hasn't been told yet. In fact, he won't learn about it for a whole week.

Later in the day, a delivery person asks Douglas, "I just saw that car parked out front. Does someone here own a Ferrari?"

Results

In order to measure judgments about knowledge and assertability, I again analyzed the results for the *combined response* the same way as in Experiment 3. A preliminary multinomial logistic regression revealed main effects of both Story and Procedure on the combined response, so I will analyze the results for each Story and Procedure separately.

Comparing the Guided conditions, there was an effect of Story on combined response, $\chi^2(6, 229) = 66.42, p < .001$. (See Fig. 4.) The magnitude of the difference in frequencies was large, Cramer's $V = .381$. Guided Control differed significantly from both Guided Non-deceptive, $\chi^2(3, 153) = 42.77, \text{Cramer's } V = .529$ (large effect size), and Guided Deceptive, $\chi^2(3, 151) = 43.79, \text{Cramer's } V = .538$ (large effect size). Guided Non-deceptive and Guided Deceptive didn't differ significantly, $\chi^2(3, 154) = 1.05, p = .79$.

In Guided Control, a majority (56%) answered that Douglas knows the proposition and should assert it, binomial test, $p < .001$, test proportion for all binomial tests reported in this experiment = .25. Only a very small number of participants (11%) answered that Douglas does not know the proposition and should assert it, binomial test, $p = .004$. In Guided Non-deceptive, a majority (62%) answered that Douglas does not know the proposition and should not assert it, binomial test, $p < .001$. A non-trivial number of participants (26%) answered that Douglas does not know the proposition and should assert it, but this didn't differ from what could be expected

by chance, binomial test, $p = .982$. In Guided Deceptive, a majority (59%) answered that Douglas does not know the proposition and should not assert it, binomial test, $p < .001$. A nontrivial number of participants (28%) answered that Douglas does not know the proposition and should assert it, but this didn't differ from what could be expected by chance, binomial test, $p = .678$.

Comparing the Unguided conditions, there was an effect of Story on combined response, $\chi^2(6, 226) = 31.36, p < .001$. The magnitude of the difference in means was medium, Cramer's $V = .263$. Unguided Control differed significantly from both Unguided Non-deceptive, $\chi^2(3, 151) = 28.68, p < .001$, Cramer's $V = .436$ (medium effect size), and Unguided Deceptive, $\chi^2(3, 149) = 18.24, p < .001$, Cramer's $V = .350$ (medium effect size). Unguided Non-deceptive and Unguided Deceptive conditions didn't differ significantly, $\chi^2(3, 152) = 3.70, p = .295$.

In Unguided Control, a majority (74%) answered that Douglas knows the proposition and should assert it, binomial test, $p < .001$. Only a very small number of participants (4%) answered that Douglas does not know the proposition and should assert it, binomial test, $p < .001$. In Unguided Non-deceptive, a plurality (38%) answered that Douglas knows the proposition and should assert it, which exceeds chance rates, $p = .018$. Only a small number of participants (17%) answered that Douglas does not know the proposition and should assert, binomial test, $p = .11$. In Unguided Deceptive, a majority of participants (51%) answered that Douglas knows the proposition and should assert it, which far exceeds chance rates, $p < .001$, and the rate at which participants selected this option in the Unguided Non-deceptive, binomial test, $p = .034$, test proportion = .38. A small number of participants (20%) answered that Douglas does not know the proposition and should assert it, binomial test, $p = .195$.

It is also worth briefly comparing the effect of Procedure for the three different Story types. In each case, the difference was significant. Guided Control differed from Unguided, $\chi^2(3, 149) = 13.65$, $p = .003$, Cramer's $V = .303$ (medium effect size). Guided Non-deceptive differed from Unguided, $\chi^2(3, 155) = 24.39$, $p < .001$, Cramer's $V = .397$ (medium effect size). Guided Deceptive differed from Unguided, $\chi^2(3, 151) = 37.71$, Cramer's $V = .500$ (large effect size). Overall, across all three story types, the basic effect of Guided procedures was to (i) diminish the rate at which participants attributed both knowledge and assertability, and (ii) increase the rate at which they denied both knowledge and assertability.

Again I analyzed how frequently participants offered a unified response for knowledge and assertability. Binary logistic regression revealed no main or interaction effects of Story or Procedure on whether participants offered a unified response; the model was statistically insignificant, $\chi^2(5, 455) = 6.70$, $p > .05$. Across the six conditions, percentage of unified responses ranged from 67% to 81%. Overall, 75% (338 of 455) of participants offered a unified response, $\chi^2(1, 455) = 105.41$.




Fig. 4. Experiment 4. Number of participants selecting the options as best describing the agent in the story. “K&S” = the agent knows the proposition and should say it. “K&~S” = the agent knows the proposition and should not say it. “~K&S” = the agent does not know the proposition and should say it. “~K&~S” = the agent does not know the proposition should not say it.

Discussion

The results again replicated the main findings from earlier experiments. Once again the pattern was consistent with the knowledge account, this time using another completely new story based on a famous thought experiment from the literature. The central tendency in all conditions was to keep knowledge and assertability together. In control conditions, this meant attribution both knowledge and assertability at very high rates. Depending on how participants were tested, in Gettier cases this meant one of two things: surprisingly high rates of attribution for both knowledge and assertability, or very high rates of denial for both.

These findings support prior conclusions on the importance of presenting cases to participants and probing them in ways that are meaningful to them (Turri, 2013b; Turri, 2014b).

The findings also support prior work which has identified a surprisingly strong tendency in folk epistemology to attribute knowledge in Gettier cases (Starmans & Friedman, 2012). Finally, the findings also suggest that adding intentional deception to Gettier cases might increase the likelihood that people will attribute knowledge to the victimized agent. With less didactic procedures in place, 60% of people attributed knowledge to a victim of intentional deception in an “apparent evidence” Gettier case. In light of this finding and prior work showing that the presence of deception can interfere with task performance, this suggests a methodological lesson. Researchers studying knowledge judgments, particularly in Gettier cases, should exclude the element of intentional deception unless their research question specifically calls for it. Or, at the very least, they should include controls to rule out low-level explanations of findings in terms of deception.

Conclusion

Researchers investigating the norm of assertion agree that their project is, at least in large part, empirical. An adequate theory “must face the linguistic data” (Douven, 2006, p. 45). Critics of the knowledge account claim that it mischaracterizes our actual practice of evaluating assertions and that it clashes with the intuitively correct verdict in particular cases, including cases of reasonable false assertions, “selfless assertions,” and “Gettiered assertions.” Prior empirical work has shown that critics were wrong about the ordinary way of evaluating reasonable false assertions and selfless assertions. This paper investigated whether the critics fared any better when it comes to cases of “Gettiered assertions.”

A Gettiered assertion supposedly has two crucial features. First, intuitively the agent does not know the relevant proposition. Second, intuitively the agent should, if asked, assert the

proposition. But when I tested several paradigm cases of such assertions, I found that the agent was consistently viewed very differently. On the one hand, I tested a classic case involving an environmental threat to the veracity of the agent's perceptual experience: a "fake barn" case (Goldman, 1976). Many philosophers assume that in such a case the agent obviously lacks knowledge even though she should, if asked, assert the relevant proposition. But an overwhelming majority of people judged that the agent both knew the proposition and should assert it. Indeed, the fake barn case was judged no differently than a closely matched control case involving cheap barns. On the other hand, I tested cases with an "explanatory disconnect" structure that more closely resembles the structure of Edmund Gettier's thought experiments (Gettier, 1963). Again, many philosophers assume that such cases obviously involve assertability without knowledge. And, again, an impressive majority of people consistently judged otherwise.

Neither sort of case tested here poses a problem in principle for the knowledge account of assertion. In particular, the results should vanquish any suspicion that fake barn cases pose even the slightest problem for the knowledge account. The results for "explanatory disconnect" cases were less overwhelmingly favorable to the knowledge account, but the central tendency was still for judgments of knowledge and assertability to remain united. And the erosion of unity I did observe could have been due to mild incredulity or confusion at what, almost inevitably, is a peculiar plot line. (The plot lines tested here are, nevertheless, far simpler and more credible than ones often found in the literature.) Whatever the explanation, the key point is that the central tendency was in line with the knowledge account and opposed to the conventional wisdom among critics.

I obtained the present results using procedures based on prior experimental work revealing

an effective way to guide participants to appreciate relevant features of a “Gettier case.” The stories were partitioned into three stages and included questions to prime participants to attend to these important details of the story (Turri 2013b). Moreover, Experiment 4’s results provide the best evidence to date that these procedures cause participants to attribute knowledge in ways consistent with the alleged consensus about Gettier cases among contemporary professional epistemologists. In particular, one group of participants was tested using the “didactic” procedures of tri-partitioning and priming questions, while another group was tested using “non-didactic” procedures that did neither of those things. Overall, the basic effect of more didactic procedures was to diminish the rate at which participants attributed both knowledge and assertability, and to increase the rate at which they denied both knowledge and assertability. Less didactic procedures produced results inconsistent with conventional wisdom about Gettier cases. In particular, rates of knowledge attribution significantly exceeded chance rates in the Gettier conditions, whereas philosophers claim that knowledge is obviously lacking in such cases. By contrast, more didactic procedures produced results consistent with conventional wisdom about Gettier cases: rates of knowledge attribution were at or near floor. The rate at which participants linked knowledge and assertability remained constant across both sorts of procedure (~75%).

Case creep and competent curation

The present findings further undermine conventional philosophical wisdom about fake barn cases and other cases with an “environmental threat” structure. Other recent studies have also found that people tend to ascribe knowledge in such cases (Colaco, Buckwalter, Stich & Machery, 2014; Turri, Buckwalter & Blouw, 2014; Turri, under review). By contrast, in study after study, people deny knowledge in cases with an “explanatory disconnect” structure

(Starmans & Friedman, 2012; Turri, 2013a; Nagel, San Juan & Mar, 2013; Turri 2013b). The lesson here is that philosophers should stop grouping into one category cases with radically different causal structures.

More generally the philosophical literature and lore on Gettier cases is a vast and confusing labyrinth built adventitiously over many decades. The nominal category “Gettier case” masks radical diversity in underlying causal structure. These differences are important in both theory and ordinary practice — so important as to render the nominal category, as we have inherited it, utterly useless. As we saw above, some “Gettier cases” elicit rates of knowledge attribution exceeding 80%, while others struggle to top 20%. The mere fact that something is a “Gettier case” is consistent with its being both overwhelmingly judged knowledge and overwhelmingly judged ignorance, thereby masking differences that radically affect the psychology of knowledge attributions and depriving the category of any diagnostic or predictive value (Blouw, Buckwalter & Turri, in press). Philosophers should not assume too much when discussing “Gettier cases.”

Philosophers should also learn a general lesson from this and beware what we might call “case creep.” Case creep occurs when there is a (near) consensus verdict about a particular original case or a few cases with essentially similar underlying structure. Call these *acme cases*. Respecting the consensus verdict about acme cases is treated as a criterion of theoretical success. Then, as the literature develops over time, researchers create new cases and add them to the discussion. These new cases resemble the acme case in important respects, but they are also importantly different in some way. After all, if the new case was essentially the same as the acme case, then there would be no point to introducing it — the acme case would work just as well. Call these *creepy cases*. Newer creepy cases are added in turn, and these resemble the acme case

even less. Each newer, creepier generation introduces more innovation and diversity. If we keep calling them all “acme cases,” then before long the common label will conceal important underlying differences, rendering the nominal category diagnostically useless. We began with a genuinely important insight about some unified original phenomenon, typified by the acme case. We end with a hodgepodge of conflicting intuitions about a creepy category without integrity, populated by a motley crew of creatures bred of momentary expedience but whose cumulative consequence is a cancer to inquiry.

How critics got it so wrong: some hypotheses

Why have critics badly mischaracterized the natural and intuitive view of Gettiered assertions? I propose the following hypotheses.

For fake barn cases, the critics’ intuitions about assertability track their implicit judgments about knowledge, courtesy of their ordinary competence in applying that concept. Knowledge attributions are social judgments and many aspects of social cognition, including moral cognition, occur automatically and unconsciously (Bargh, Schwader, Hailey, Dyer & Boothby, 2012; Haidt, 2007). It seems clear that the agent should make the assertion because she clearly knows. But contemporary philosophers have also been trained to say that someone in “fake barn country” obviously lacks knowledge. So the default knowledge judgment is explicitly reversed. They do not realize that they are explicitly denying the intuitive basis for attributing assertability. This causes them to misinterpret such cases as problematic for the knowledge account, even though the cases actually support it.

For other types of Gettier case, especially those featuring intentional deception or where the agent isn’t fundamentally misled about the nature of his evidence, the explanation could be

slightly more complicated. I see two paths to the mistaken official view of critics. On the one hand, in virtue of their folk competence in applying the concept of knowledge, many philosophers feel the attraction of attributing knowledge in such cases. This implicitly leads them to judge that the agent should make the assertion. They are not trained to think carefully about assertability in Gettier cases, so this default attribution of assertability stands. But they are trained to think carefully about and, furthermore, to deny knowledge in Gettier cases. So the default knowledge judgment is explicitly reversed. They do not realize that they are explicitly denying the intuitive basis for attributing assertability. This causes them to misinterpret these cases as problematic for the knowledge account, even though the cases actually support it. This explanation is parallel to the explanation of mistaken professional opinion about fake barn cases discussed above.

On the other hand, I repeatedly observed that roughly 1-in-5 people attribute knowledge but deny assertability in “explanatory disconnect” Gettier cases. This did not exceed chance rates, but it arguably suggests that there are important individual differences in people’s willingness to adopt this combination of attitudes. Future work could investigate what these differences are and whether philosophers tend to have traits that predict greater willingness to adopt the combination, such as styles of reasoning and cognition (Epstein, Pacini, Denis-Raj & Heier, 1996; Stanovich & West, 1998; Kosslyn & Miller, 2013).

Whatever the ultimate explanation, philosophers should reconsider their uncritical reliance on conventional wisdom about Gettier cases and, I suspect, knowledge judgments more generally.

Coda

In conclusion, critics have claimed that the knowledge account is counterintuitive and mischaracterizes our ordinary practice of evaluating assertions. To substantiate these charges, critics have used thought experiments to pump intuitions against the knowledge account, including cases involving reasonable false assertions, selfless assertions, and Gettiered assertions. Critics claim that, intuitively, these are cases of assertability without knowledge. Prior work has shown that critics were wrong about reasonable false assertions and selfless assertions. In this paper, we saw that critics are also wrong about Gettiered assertions, which pose no problem in principle for the knowledge account. Critics have hit an ignominious intuitional trifecta. Combine this with the fact that the knowledge account is supported by a wealth of positive theoretical and empirical evidence (Turri 2016), and we've reached a watershed in research on the norm of assertion. It's time to move past the hunt for alleged counterexamples which supposedly show that knowledge isn't the norm of assertion after all. Further progress in this area will come by trying to better understand the precise sense in which knowledge is the norm of assertion and, ultimately, why knowledge specifically plays this role.

Acknowledgements — For feedback I thank Wesley Buckwalter, David Rose, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

References

Bach, K., & Harnish, R. M. (1979). *Linguistic communication and speech acts*. Cambridge,

Mass. : MIT Press.

- Bargh, J. A., Schwader, K. L., Hailey, S. E., Dyer, R. L., & Boothby, E. J. (2012). Automaticity in social-cognitive processes. *Trends in Cognitive Sciences*, 16(12), 593–605. doi:10.1016/j.tics.2012.10.002
- Blouw, P., Buckwalter, W., & Turri, J. (2015). Gettier cases: a taxonomy. In R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining knowledge: new essays on the Gettier problem*. Oxford: Oxford University Press.
- Chisholm, R. (1989). *Theory of Knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Coffman, E. J. (2014). Lenient accounts of warranted assertability. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: new essays on action, belief and assertion* (pp. 33–59). Oxford University press.
- Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(02), 199–212. doi:10.1017/epi.2014.7.
- Cullen, S. (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1(2), 275–296. doi:10.1007/s13164-009-0016-1
- Douven, I. (2006). Assertion, Knowledge, and Rational Credibility. *Philosophical Review*, 115(4), 449–485. doi:10.1215/00318108-2006-010
- Epstein, S., Pacini, R., Denes-Raj, V., & Heier, H. (1996). Individual differences in intuitive–experiential and analytical–rational thinking styles. *Journal of Personality and Social Psychology*, 71(2), 390. doi:10.1037/0022-3514.71.2.390
- Garson, G. D. (2012). *Logistic regression: binary and multinomial*. Asheboro, NC: Statistical

Associates Publishers.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.

Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel.

Goldman, A. I. (2009). Williamson on knowledge and evidence. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 73–91). Oxford: Oxford University Press.

Haidt, J. (2007). The New Synthesis in Moral Psychology. *Science*, 316(5827), 998–1002.
doi:10.1126/science.1137651

Hetherington, S. (2013). There can be lucky knowledge. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology*. (2nd ed.). Malden, Mass.: Wiley-Blackwell.

Hill, C., & Schechter, J. (2007). Hawthorne's lottery puzzle and the nature of belief. *Philosophical Issues*, 17, 102–122.

Kosslyn, S. M., & Miller, W. G. (2013). *Top brain, bottom brain*. New York: Simon & Schuster.

Kvanvig, J. (2008). Epistemic Luck. *Philosophy and Phenomenological Research*, 77(1), 272–281. doi:10.1111/j.1933-1592.2008.00187.x

Kvanvig, J. (2009). Assertion, knowledge, and lotteries. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge* (pp. 140–160). Oxford: Oxford University Press.

Lackey, J. (2007). Norms of assertion. *Nous*, 41(4), 594–626.

Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25(5), 168–175.

Lycan, W. G. (2006). The Gettier problem problem. In S. Hetherington (Ed.), *Epistemology futures* (pp. 148–168). Oxford: Oxford University Press.

Matilal, B. K. (1986). *Perception: an essay on classical Indian theories of knowledge*. Oxford:

Oxford University Press.

- Nagel, J., San Juan, V., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 1–10. doi:10.1016/j.cognition.2013.02.008
- Neta, R., & Rohrbaugh, G. (2004). Luminosity and the Safety of Knowledge. *Pacific Philosophical Quarterly*, 85(4), 396–406.
- Pritchard, D. (2014). Epistemic luck, safety, and assertion. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: new essays on action, belief and assertion*. Oxford: Oxford University Press.
- Reynolds, S. L. (2002). Testimony, knowledge, and epistemic goals. *Philosophical Studies*, 110(2), 139–161.
- Sartwell, C. (1992). Why knowledge is merely true belief. *The Journal of Philosophy*, 89(4), 167–180.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 85(2), 383–403.
- Schul, Y., Mayo, R., Burnstein, E., & Yahalom, N. (2007). How people cope with uncertainty due to chance or deception. *Journal of Experimental Social Psychology*, 43(1), 91–103. doi:10.1016/j.jesp.2006.02.015
- Smithies, D. (2012). The Normative Role of Knowledge. *Nous*, 46(2), 265–288. doi:10.1111/j.1468-0068.2010.00787.x
- Sperber, D., & Noveck, I. A. (2004). Introduction. In D. Sperber & I. A. Noveck (Eds.), *Experimental pragmatics* (pp. 1–22). New York: Palgrave Macmillan.
- Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of*

- Experimental Psychology: General, 127(2), 161. doi:10.1037/0096-3445.127.2.161
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272–283. doi:10.1016/j.cognition.2012.05.017
- Turri, J. (2010a). Prompting challenges. *Analysis*, 70(3), 456–462. <http://doi.org/10.1093/analys/anq027>
- Turri, J. (2010b). Epistemic invariantism and speech act contextualism. *Philosophical Review*, 119(1), 77–95. <http://doi.org/10.1215/00318108-2009-026>
- Turri, J. (2012). Is knowledge justified true belief? *Synthese*, 184(3), 247–259. <http://doi.org/10.1007/s11229-010-9773-8>
- Turri, J. (2013a). A conspicuous art: putting Gettier to the test. *Philosophers' Imprint*, 13(10), 1–16.
- Turri, J. (2013b). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, 129(2), 279–291. <http://doi.org/10.1016/j.cognition.2013.06.012>
- Turri, J. (2014). The problem of ESEE knowledge. *Ergo*, 1(4), 101–127.
- Turri, J. (2015a). Knowledge and the norm of assertion: a simple test. *Synthese*, 192(2), 385–392. <http://doi.org/10.1007/s11229-014-0573-4>
- Turri, J. (2015b). Selfless assertions: some empirical evidence. *Synthese*, 192(4), 1221–1223. <http://doi.org/10.1007/s11229-014-0621-0>
- Turri, J. (2016). Knowledge and the norm of assertion: an essay in philosophical science. Open Book Publishers.
- Turri, J. (in press). Knowledge, certainty and assertion. *Philosophical Psychology*.
- Turri, J. (under review). Vision, knowledge, and assertion. University of Waterloo.

- Turri, J., & Blouw, P. (2015). Excuse validation: a study in rule-breaking. *Philosophical Studies*, 172(3), 615–634. <http://doi.org/10.1007/s11098-014-0322-z>
- Turri, J., Buckwalter, W., & Blouw, P. (2015). Knowledge and luck. *Psychonomic Bulletin & Review*, 22(2), 378–390. <http://doi.org/10.3758/s13423-014-0683-5>
- Unger, P. (1975). *Ignorance: a case for skepticism*. Oxford: Oxford University Press.
- Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1&2), 429–460.
- Weiner, M. (2005). Must we know what we say? *Philosophical Review*, 114(2), 227–251.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105(4), 489–523.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, S. (2014). The dual-aspect norms of belief and assertion. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: new essays on action, belief and assertion* (pp. 239–258). Oxford University Press.