Experimental Epistemology and "Gettier" Cases*

John Turri
john.turri@gmail.com

Abstract: This chapter reviews some faults of the theoretical literature and findings from the experimental literature on "Gettier" cases. Some "Gettier" cases are so poorly constructed that they are unsuitable for serious study. Some longstanding assumptions about how people tend to judge "Gettier" cases are false. Some "Gettier" cases are judged similarly to paradigmatic ignorance, whereas others are judged similarly to paradigmatic knowledge, rendering it a theoretically useless category. Experimental procedures can affect how people judge "Gettier" cases. Some important central tendencies in judging "Gettier" cases appear to be robust against demographic variation in biological sex, age, language, and culture, although there could be some interesting differences related to culture and personality traits. Some remaining questions regarding Gettier's cases, "Gettier" cases, and "the Gettier problem" concern the psychology and sociology of contemporary anglophone theoretical epistemology. Some remaining questions regarding the empirical study of knowledge judgments concern mechanisms underlying observed behavioral patterns.

keywords: experimental epistemology; philosophical methodology; thought-experimenter bias

word count: 7113

---

* This is the penultimate version (October 2017) of a paper to appear in S. Hetherington (Ed.), *The Gettier problem* (Cambridge University Press). Please cite the final, published version if possible.

Gettier's Cases

Edmund Gettier's 1963 paper, "Is Justified True Belief Knowledge?", was explicitly concerned with evaluating "various attempts" that had "been made in recent years to state necessary and sufficient conditions for someone's knowing a given proposition." Gettier claimed that those "attempts have often been such that they can be stated in a form similar to" a definition stating that *you know that P* just in case *you have a justified true belief that P*. Gettier claimed, furthermore, that the latter was insufficient for the former, providing two examples to support his claim. Here is one of Gettier's original examples (Gettier 1963: 122-123, "Case II"):

> Let us suppose that Smith has strong evidence for the following proposition:
>
> (f) Jones owns a Ford.
>
> Smith's evidence might be that Jones has at all times in the past within Smith's memory owned a car, and always a Ford, and that Jones has just offered Smith a ride while driving a Ford. Let us imagine, now, that Smith has another friend, Brown, of whose whereabouts he is totally ignorant. Smith selects three place-names quite at random, and constructs the following three propositions:
>
> (g) Either Jones owns a Ford, or Brown is in Boston;
>
> (h) Either Jones owns a Ford, or Brown is in Barcelona;
>
> (i) Either Jones owns a Ford, or Brown is in Brest-Litovsk.

> Each of these propositions is entailed by (f). Imagine that Smith realizes the entailment of each of these propositions he has constructed by (f), and proceeds to accept (g), (h), and (i) on the basis of (f). Smith has correctly inferred (g), (h), and (i) from a proposition for which he has strong evidence. Smith is therefore completely justified in believing each of these three propositions. Smith, of course, has no idea where Brown is.
>
> But imagine now that two further conditions hold. First, Jones does *not* own a Ford, but is at present driving a rented car. And secondly, by the sheerest coincidence and entirely unknown to Smith the place mentioned in proposition (h) happens really to be the place where Brown is. If these two conditions hold then Smith does *not* know that (h) is true, even though (*i*) (h) *is* true, (*ii*) Smith does belief that (h) is true, and (*iii*) Smith is justified in believing that (h) is true.

One leading philosopher labeled Gettier's paper "Gettier's survey" (Jackson 2011: 480–481), which "invited" philosophers "to agree with his intuition" that the examples were not cases of knowledge (Jackson 1998: 28). Evidence from surveys on whether people attribute or deny knowledge in specific cases is "highly relevant to what their concept of knowledge is," so the contribution of Gettier's paper was to provide "empirical, a posteriori" evidence that "so many readers agreed with Gettier that the cases he presented were not cases of knowledge" (Jackson 2011: 476-477). Literature reviews claim that "most" philosophers who consider such cases agree with Gettier (e.g. Turri 2012a: 215; Hetherington 2016: ix)

Several critical points should be made immediately about such claims and Gettier's cases. First, it is not actually known whether "many" or "most" philosophers agree about such cases, because no empirical study of philosophers' judgments of Gettier's cases has ever been published. Accordingly, any claim about the relevant proportions is speculative and ought to be treated as such. Second, the description of Gettier's paper as a "survey," even in an attenuated and extended sense, is inaccurate and misleading. Gettier did not "invite" philosophers to share his intuition. Instead he used language apt to prime attributions of ignorance: Smith is described as selecting propositions "at random," accepting them despite having "no idea" regarding principal facts implicated by the propositions, and also ending up with a true belief "by the sheerest coincidence" for reasons "entirely unknown" to him. Gettier did not probe for judgments. Instead he confidently and unqualifiedly inserted his own verdicts into the description of the case. Despite the unnaturalness and pointlessness of Smith's inferences, Gettier tells us that Smith is "completely justified" in making them, and Gettier tells us that "Smith does not know" the relevant proposition despite having a justified true belief. Third, and relatedly, when judged as survey instruments intended to probe for knowledge judgments, Gettier's cases are miserably constructed. In addition to being "contrived and artificial" (Dancy 1985: 26), they are long, stilted, tendentious, multiply confounded, and not paired with relevant controls.

For these reasons, even if one tested Gettier's cases and found that people unanimously judged "Smith doesn't know," it would provide little if any information about "our folk theory" of knowledge, contrary to what some have suggested (e.g. Jackson 1998: 31). Similarly, if one thought that the concept of knowledge at issue in the contemporary literature was "a philoso-

phers' artifact" unconnected to "anything possessed by ordinary people" (Lycan 2006: 165), testing Gettier's cases would provide negligible information about the standards implicit in professional philosophers' judgments. In short, Gettier's cases are completely unsuitable for any serious or worthwhile attempt to gain evidence regarding people's knowledge concept. In particular, they do not "prov[e] that justified true belief is insufficient for knowledge" (Jacquette 1996: 115-116), according to any knowledge concept.

As if this were not bad enough, adding to this embarrassment of contemporary epistemology, much simpler cases had long existed, which are better suited to teach the lesson supposedly learned from Gettier's cases, as prominent philosophers have pointed out for decades (e.g. Goldman 1967: 357, fn. 1; Matilal 1986: 135-137; Chisholm 1989: 92-93). Moreover, there is no evidence that the view targeted by Gettier's cases — the "justified true belief" or "JTB" theory of knowledge — was ever held by historically influential philosophers (Kaplan 1985; Plantinga 1993: 6-7; Dutant 2015), let alone any evidence for the wildly irresponsible claim that it was "the most widely accepted definition of (propositional) knowledge in the history of philosophy" (Jacquette 1996: 115).

The serious defects undermining Gettier's original cases are not limited to those particular cases. Another influential case from early in the literature also focused on car ownership (Lehrer 1965; 169–170). The agent in this example, Keith, has an "honest and reliable" friend named "Mr. Nogot." Nogot gets out of a new Ford, walks into Keith's office, tells Keith that he has "just purchased the car," and – weirdly – "shows [Keith] a certificate that states that he owns the Ford." On this basis, Keith believes, reasonably, "Mr. Nogot, who is in my office, owns a Ford."

Then Keith deduces, “Someone in my office owns a Ford.” We are told that Keith is “completely justified” in making this inference, despite its unnaturalness and apparent pointlessness. We are then told that Mr. Nogot has “deceived” Keith and does not own a Ford. We are left to guess why Mr. Nogot, honest and reliable friend that he is, would do this. Nevertheless, Keith also sees another person in his office, “Mr. Havit,” who does own a Ford. But Keith has “no evidence that [Mr. Havit] owns a Ford.” Thus it is true that someone in Keith’s office owns a Ford. However, we are told, Keith does “not know that it is true.” Again, a weirdly contrived scenario is described using tendentious language and all of the relevant verdicts are confidently inserted into the case’s description.

Here is a another, recent example of a “Gettier” case offered in apparent seriousness. Although, it might be more accurately labeled a “chain linked Gettier” case because it consists of two parts, one which explains how an agent is “Gettiered” and another which then attempts to embed that into a more complicated situation:

> (*Expert botanist*) David is an expert botanist, able to competently distinguish between the over 20,000 different species of orchid. David is presented with an orchid and asked to identify its species. Using his amazing skill, he can clearly tell that this particular orchid is either going to be a *Platanthera tescamnis* or a *Platanthera sparsiflora* (which look quite similar), and upon even further expert analysis he comes to the conclusion that it is a *Platanthera tescamnis* orchid, which it is. However, Kevin, David’s nemesis and an expert botanist in his own right, decided the night before to disguise the *Platanthera tescamnis* orchid to look like a *Platanthera sparsiflora* or-

chid. Thankfully, however, Alvin, David's other expert botanist nemesis (who is conveniently not on speaking terms with Kevin), decided to try to trick David in the same way— arriving shortly after Kevin left, perceiving that the orchid was a *Platanthera sparsiflora*, and disguising Kevin's disguise to look like a *Platanthera tescamnis*, which it happens to actually be. (Church 2013: 174)

(*Orchid guessing*) Ruth, Dave, Shelly, and Bob are playing a guess-the-species-of-orchid game — a game where they are presented with various types of orchid and asked to identify the species. Ruth is an expert botanist. When she is presented with an orchid she is able to use her immense skills as a botanist to narrow down the over 20,000 possibilities to the one right answer. Dave too is an expert botanist; indeed, he is every bit as knowledgeable as Ruth. Using his amazing skill he can clearly tell that the orchid before him is either going to be a *Platanthera tescamnis* or a *Platanthera sparsiflora* (which look quite similar), and upon even further expert analysis he comes to the conclusion that it is a *Platanthera tescamnis* orchid, which it is. However, after narrowing down the over 20,000 possibilities to just *Platanthera tescamnis* and *Platanthera sparsiflora*, Dave is Gettiered about these final two options. Thirdly, Shelly is presented with an orchid. She knows almost nothing about botany; however, she is something of an idiot savant — having memorized the names of every single species of orchid. She has no idea what species is before her, so she simply picks a species at random and just happens to get it right. Finally, Bob is presented with an orchid. He knows absolutely nothing about botany. He doesn't even know the names of any

> species. When he goes to hazard a guess; however, he chokes on a burp and just so happens to utter the name of the species before him. (Church 2013: 175-176)

Afterward — in the main text, after presenting the cases and while stipulating verdicts about them — the author of these examples clarifies that Dave's "circumstances are just like David's in the case of Expert Botanist" (Church 2013: 176). Again we are confronted with a long, complicated description of an exceedingly contrived situation. (Or is it two situations?) Despite the length, some critical details are omitted and later stipulated elsewhere. Tendentious language abounds — "knows almost nothing," "is Gettiered" (!), "has no idea," "guessing," "picks [an answer] at random and just happens to get it right," "knows absolutely nothing," "goes to hazard a guess," "just so happens to utter the name of the species." I leave it to others to speculate on the significance of the fact that some professional philosophers believe that we advance our understanding of anything by considering whether an agent's intellectual performance is "on a par with Bob's choked-on burp" (Church 2013: 176).

Arguably we can detect an underlying structure in some of these poorly constructed cases (Zagzebski 1996; Jacquette 1996; Feldman 2003). It involves a person with imperfect but potentially good evidence for thinking that a specific proposition is true. This person notices that this first proposition entails a second proposition, which he concludes is true. And it is true. However, the impressive evidence he began with turns out to be misleading and he never detected that the first proposition is true. In fact, the first proposition is false. In a proper study of judgments about the case, we would now decide two things. On the one hand, does the person have a justified true belief that the second proposition is true? On the other hand, does the person know that the sec-

ond proposition is true? If the central tendency is for our answers to be "yes" and "no," respectively, then we are committed to denying that knowledge is equivalent to justified true belief.

Researchers have recently made serious, unbiased attempts to test cases structured this way, to which we turn next. In what follows, I distinguish between *Gettier's cases*, which I have already discussed, and *"Gettier" cases* (always with scare quotes). The former come from Gettier's original paper, but the latter are merely a nominal category with no underlying unity (Turri 2016a). For reasons already discussed, no serious study of knowledge judgments would use Gettier's cases because the results could not be meaningfully interpreted. As the theoretical literature developed over the decades, a bewildering array of cases have been carelessly labeled "Gettier" cases. Based on the disappointingly uncritical reception of Gettier's cases, it is perhaps not surprising that philosophers have been injudicious in keeping track of whether and how new "Gettier" cases differ from Gettier's cases or each other, and whether such differences are important. By contrast, in just a few years, the experimental research discussed below has made significant progress in identifying factors affecting central tendencies in knowledge judgments. Of course, once a case's structure is clarified, serviceable stimuli are constructed, and appropriate controls are included, the study has potential to be informative regarding knowledge judgments. Whether we call it a "Gettier" case is immaterial.

Experimental Research on "Gettier" Cases

The first study of "Gettier" cases I am aware of was motivated by psychological research sug-

gesting important cultural differences in reasoning styles and moral judgments (Weinberg, Nichols & Stich 2001). Researchers tested a story about car ownership on U.S. undergraduates from a variety of cultural backgrounds. The story was similar to some cases from the early Gettier literature, but it was phrased much more naturally and unbiasedly.

> Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it?

Approximately 25% of those reporting Western cultural backgrounds attributed knowledge but more than half of those reporting eastern or southern Asian backgrounds attributed knowledge. The results fit nicely with prior cross-cultural work on other sorts of judgment. Followup work on knowledge judgments has not consistently replicated these cultural differences when using the same materials on undergraduates in the United States (Kim and Yuan 2015), when using slightly modified materials on lay populations in the United Kingdom (Seyedsamymodst 2015), when using modified materials and procedures to test residents of the U.S. and India (Turri 2013: section 7), or, on some ways of probing for knowledge judgments, when using different materials to test residents of Brazil, India, Japan, and the United States in their native languages (Machery et al. 2015). This followup work has observed consistently low rates of knowledge attribution, typically around 20%. Some of this work included closely matched control conditions where people attributed knowledge.

Nevertheless, at least three separate studies have found a statistically significant difference whereby people with Asian cultural backgrounds attribute knowledge at higher rates than "Westerners" do in such cases (Weinberg, Nichols & Stich 2001; Kim & Yuan 2015, p. 356, f. 3, unfortunately relegating this finding to a footnote in a paper entitled, "No cross-cultural differences in Gettier car case intuition"; Machery et al. 2015). Another cross-cultural study, by far the most wide-ranging one on knowledge attribution in "Gettier" cases to date, found substantial agreement in knowledge judgments across 17 languages, with one group, Israeli Bedouins, diverging radically from the rest. So, at this point, the balance of evidence supports the conclusion that there is a cross-culturally robust central tendency to deny knowledge in some "Gettier" cases on some ways of probing. At the same time, there is also some of evidence that other, seemingly simpler ways of probing erase and potentially even reverse that tendency, and there is a nontrivial chance that some cultural variation exists in these matters.

At present, the evidence also suggests that participants' biological sex and age have little if any effect on knowledge judgments about "Gettier" cases (e.g. Starmans & Friedman 2012; Turri 2013; Kim & Yuan 2015; Machery et al. 2017; Turri 2017). Recent evidence suggests that individual differences, such as personality traits, might affect such judgments. A large cross-cultural study found a negative correlation between conscientiousness and denying knowledge in a "Gettier" case (Machery et al. 2017). That is, the more conscientious a participant was, the more likely the participant was to attribute knowledge in a "Gettier" case. By contrast, neuroticism and openness to new experience positively correlated with denying knowledge. These tantalizing results are likely to spark debate and speculation, which is natural and appropriate, but, as the re-

searchers themselves admirably emphasize, further research is needed before drawing confident conclusions about these matters. In other words, theory should not get too far ahead of the data.

Researchers have investigated whether questioning procedures affect the rate of knowledge attribution in "Gettier" cases. For example, participants are less likely to attribute knowledge in a "Gettier" case when the scenario is broken up into multiple parts on three separate screens, rather than being presented all at once on a single screen, and when participants are asked questions that make certain features of the case salient (Turri 2013; Turri 2016b: experiment 4). At present, it is unclear why this happens. To take another example, when participants read the case about Bob's car (see above), researchers observed a significantly lower rate of knowledge attribution when the response options were "really knows/only believes" than when they were "knows/does not now" (Cullen 2010: 288; Turri 2016: experiment 4). Other researchers first probed for knowledge attributions with the options "Yes, he knows/No, he doesn't know," followed by (within-subjects) "He knew/He thought he knew but he did not actually know" (Machery et al. 2015; Machery et al. 2017; compare Nagel, Mar & San Juan 2013). Knowledge attribution was much higher for the "Yes/No" option, a pattern that is cross-culturally robust. Interpreting this last finding is complicated, however, by the fact that the verbal difference between the pairs of options was confounded with order, length, and complexity, as well as other intervening questions about the case. For instance, maybe the mere fact that people are being asked about knowledge *again* leads them to attribute it less, or maybe they are less likely to answer affirmatively for longer or more complex response options. Nevertheless, overall the evidence clearly shows that questioning procedures can affect the psychological processing involved in producing or

recording knowledge judgments, perhaps by triggering slightly different knowledge concepts, by causing the same concept to be construed in different ways, or by changing the way participants interpret the task. Further research is required to investigate the matter.

Arguably the most important study of "Gettier" cases to date distinguished between *apparent* and *authentic* evidence (Starmans & Friedman 2012). Apparent evidence is "evidence that appears to be informative about reality, but is not really," whereas authentic evidence is, roughly, evidence that makes the belief true when based on it (Starmans & Friedman 2012: 280). To illustrate the difference, consider two versions of a story about Corey, who collects coins in his piggy bank. One day Corey looks at a quarter he is putting into his bank and notices that it looks pretty old. He checks the date and reads "1936." In the authentic-evidence version of the story, the coin is from 1936. In the apparent-evidence version it is from 1938 and part of the date has rubbed off. In each version there is already a 1936 quarter buried deep in the piggy bank, but Corey isn't aware of this other quarter. Then Corey takes a short nap, during which his roommate comes home, takes the quarter that Corey just deposited in the bank, and leaves. Corey wakes up soon after and does not realize what his roommate did. Here is the complete text of both versions:

> *(Apparent evidence*) Corey has been collecting coins in his piggy bank for years. One day he is about to put a quarter in his piggy bank, and notices that it looks pretty old. Though he's never paid attention to dates before, he reads the date and sees that it's from 1936. However, he doesn't realize that the date has partially rubbed off and it is really from 1938. There is already a quarter dated 1936 buried deep in his piggy bank, but Corey isn't aware of this. He deposits the quarter and goes to take a nap.

Corey's roommate Scott comes home, and needs some change for the bus. He shakes the piggy bank and the quarter Corey just put in falls out. Scott takes it and leaves. Corey wakes up after a 10-minute nap, and doesn't realize that Scott was there.

(*Authentic evidence*) Corey has been collecting coins in his piggy bank for years. One day he is about to put a quarter in his piggy bank, and notices that it looks pretty old. Though he's never paid attention to dates before, he reads the date and sees that it's from 1936. However, he doesn't realize that 1936 is the year his grandmother was born. There is already a quarter dated 1936 buried deep in his piggy bank, but Corey isn't aware of this. He deposits the quarter and goes to take a nap. Corey's roommate Scott comes home, and needs some change for the bus. He shakes the piggy bank and the quarter Corey just put in falls out. Scott takes it and leaves. Corey wakes up after a 10 minute nap, and doesn't realize that Scott was there.

When Corey wakes up from his nap, does he "really know" or does he "only believe" that there is a 1936 coin in his piggy bank? People who read the authentic-evidence version tended to attribute knowledge, but people who read the apparent-evidence version tended to deny knowledge. The basic finding that people tend to deny knowledge in apparent evidence cases has been replicated (e.g. Turri 2013, section 2), and so has the finding that people tend to attribute knowledge exceeding chance rates in authentic evidence cases (e.g. Nagel, San Juan & Mar 2013, using very different questioning procedures; for important discussion and corrections regarding the statistical analyses, see Starmans & Friedman 2013: 664; see also Powell, Horne, Pinillos & Holyoak 2013, using a very different measure involving false recall).

Building on the apparent/authentic distinction, more recent work has shown that the structure of "Gettier" cases differs in at least three important ways (Turri, Buckwalter & Blouw 2014; Turri, Buckwalter & Blouw 2015). First, many differ in whether the agent *initially perceives* a state of affairs that makes his or her belief true (a "truth-maker," for short). In some examples, the agent perceives a truth-maker, but in others the agent perceives a convincing fake or something which seems to entail that the relevant proposition is true. Second, many examples differ in whether the agent's perceptual relation *remains intact* throughout. Sometimes the agent perceives a certain truth-maker and events threaten to disrupt that truth-maker, but the threat ultimately fails. Other times the threat succeeds in disrupting the original truth-maker, which is then replaced by a "backup" truth-maker. Third, many examples differ in *how similar* the perceived truth-maker and backup truth-maker are. Sometimes they very closely resemble one another, while other times they differ greatly.

Current evidence suggests that all three differences affect knowledge judgments (Turri, Buckwalter & Blouw 2015). In one study, participants read one of seven versions of a story. One version was a "knowledge control," intended to elicit very high rates of knowledge attribution. Another version was an "ignorance control," intended to elicit very low rates of knowledge attribution. The other five versions combined different permutations of the three structural variables noted above. The basic storyline featured an agent, Emma, admiring jewelry in a fancy department store. Emma purchases a stone from the diamond display, puts it in her pocket, browses for another minute, then leaves the store. The different versions of the story vary whether the stone is a real diamond or a fake, whether there is a threat to the stone remaining in Emma's pocket,

whether the threat fails or succeeds, and whether any other stone also ends up in Emma's pocket. In the terminology introduced above, the different versions manipulated whether Emma *detects* an initial truth-maker for her belief that there is a diamond in her pocket as she leaves the store, whether Emma's truth-detection is saliently *threatened*, whether the threat *succeeds* in disrupting the initial truth-maker, and whether the backup truth-maker is highly *similar or dissimilar* to the initial.

In all versions, Emma purchases a stone from a jewelry store, puts it in her pocket, and soon walks out of the store. In all the stories that involve detection, the stone she purchases is a diamond. In all the stories that do not involve detection, the stone is a fake. In all the stories that involve similar backup truth-makers, the backup truth-maker is that, one way or another, a diamond is put into Emma's pocket before she leaves the store. In all the stories that involve dissimilar backup truth-makers, the backup truth-maker is that a real diamond was secretly sewn into Emma's pocket by a previous owner long ago. Table 1 summarizes the seven versions of the story.

Table 1. Description of the seven versions of the story (from Turri, Buckwalter & Blouw 2015).

Condition	*Description*
1. Knowledge Control	The stone Emma purchases is a diamond. She walks out of the store and nothing else happens.
2. Failed Threat	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, but he fails.
3. Detection Similar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
4. Detection Dissimilar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
5. No Detection Similar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
6. No Detection Dissimilar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
7. Ignorance Control	The stone Emma purchases is a fake. She walks out of the store and nothing else happens.

By one estimation or another, stories 2-6 would all count as "Gettier" cases (e.g. Goldman 1976; Sosa 1991, 238-239; Zagzebski 1996, 283-285; Pritchard 2005; Greco 2010: chapter 5). Story 2 is structurally similar to the "fake barn" case, and story 6 is structurally similar in some important ways to Gettier's original cases. But people judged the stories very differently (see Figure 1). Rates of knowledge attribution for story 2 were very high (over 80%) and did not dif-

fer from rates observed for the knowledge control story. By contrast, rates of knowledge for story 6 were extremely low (under 20%) and did not differ from rates observed for the ignorance control story. Rates for the other three stories fell somewhere in between. Researchers have replicated this same basic pattern of results using different cover stories and procedures.

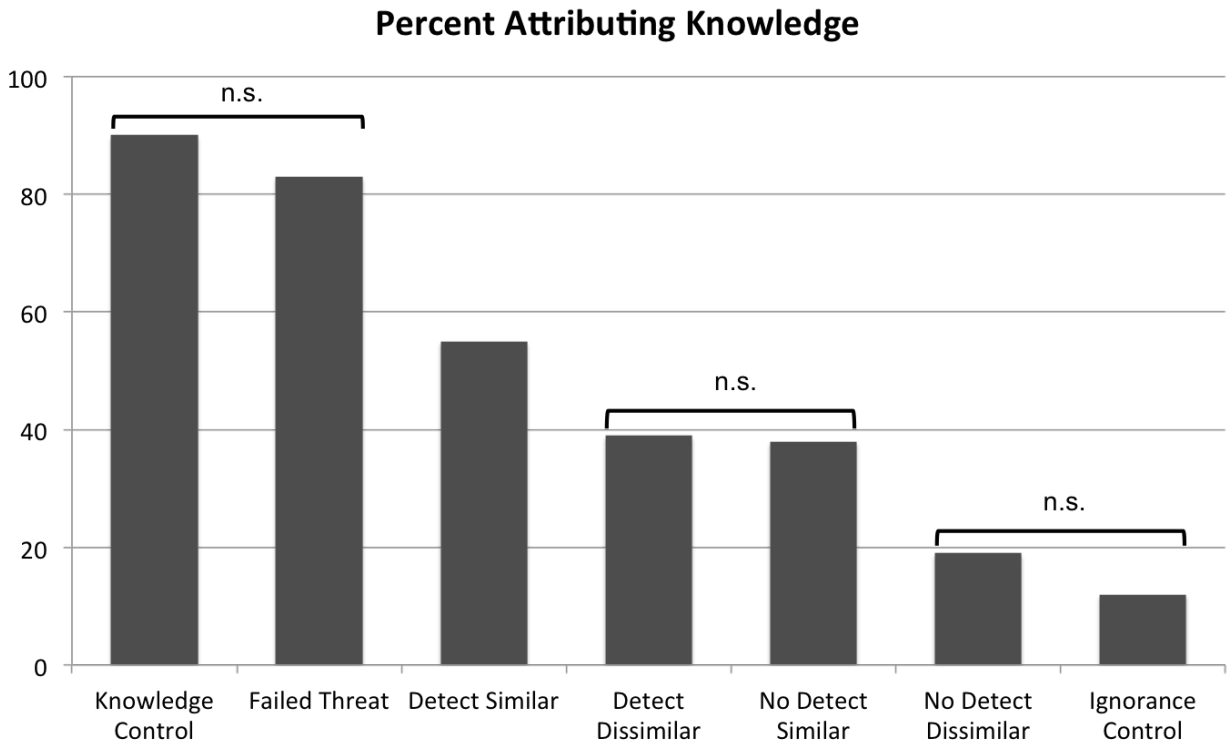

Fig. 1. Percent of participants attributing knowledge across conditions. Except where non-significance is indicated, significance for all comparisons at the $p < 0.01$ level (from Turri, Buckwalter & Blouw 2015).

Of course, there is no reason to expect that the taxonomy identifies all the factors that affect knowledge judgments or all the criteria implicit in our ordinary knowledge concept. The re-

searchers who proposed the taxonomy never made such ambitious and unsubstantiated claims, and existing research on folk epistemology clearly demonstrates that other factors also matter, including a qualitative difference between knowing positively and negatively valanced outcomes (Beebe & Buckwalter 2010; Buckwalter 2013; Beebe & Jensen 2013; Turri 2014), a qualitative difference between probabilistic information pertaining to propensities and base-rates (Turri & Friedman 2014; Friedman & Turri 2015; Turri in press), a qualitative difference between different informational sources (Turri 2015a, Turri 2015b), a qualitative difference between knowing affirmations and negations (Turri 2017a; Turri ms), and the relationship between knowledge and actionability (Turri & Buckwalter 2017; Turri, Buckwalter & Rose 2016). Researchers are fully aware that further work remains to be done (Blouw, Buckwalter & Turri in press). Instead, the value and purpose of the taxonomy is that it systematically identifies several factors that affect knowledge judgments, whose significance the theoretical literature has failed to consistently and explicitly reflect, and which, at this point, will inform any responsible treatment of "Gettier" cases. Hopefully future research advances to the point where the taxonomy is surpassed and no longer generates further research questions in this area.

In the meantime, the taxonomy just discussed can be used to catalog and even predict results from other studies of "Gettier" cases. For example, stories 3 ("detect similar") and 5 ("no detect similar") are structurally similar to the "authentic evidence" and "apparent evidence" scenarios, respectively, previously tested by researchers. Based on how the taxonomy classifies these cases, the prediction is that, other things being equal, people will tend to attribute knowledge in authentic-evidence cases and to deny knowledge in apparent-evidence cases. As reported

above, this is exactly what researchers observed. Also, earlier I mentioned that story 2 ("failed threat") is structurally similar to "fake barn" cases, which researchers had already begun studying. In the epistemology literature, theorists have claimed, without providing evidence, that "we would be strongly inclined" to not attribute knowledge in such cases (Goldman 1976; see also Sosa 1991: 238–239; Neta & Rohrbaugh 2004: 401; Pritchard 2005: 161–162; Kvanvig 2008: 274). And philosophers have relied on this verdict in order to evaluate or motivate theories of knowledge. However, based on how the taxonomy classifies this case, the prediction is that, other things being equal, people will tend to attribute knowledge.

The first study of judgments regarding a "fake barn" case used this scenario:

> Suzy looks out the window of her car and sees a barn near the road, and so she comes to believe that there's a barn near the road. However, Suzy 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 fake barn facades in this area that look as though they are real barns. In fact, Suzy is looking at the only real barn in the area. (Swain, Alexander & Weinberg 2008: 154–155)

Participants then rated the statement, "Suzy knows there is a barn near the road," on a 5-point Likert scale ("strongly disagree" = 1, through "strongly agree" = 5). Researchers found that participants tended to attribute knowledge, with the mean rating (= 3.6) exceeding the neutral midpoint (= 3) and came close to the mean rating observed for what the researchers judged to be a "clear case of knowledge" (= 3.9) (Swain, Alexander & Weinberg: 143, 146). However, the "clear case of knowledge," involving a chemist, was very different from the "fake barn" case, so

interpreting this comparison is difficult. Another study tested the same "fake barn" case involving Suzy but used a dichotomous measure for the knowledge attribution. Again researchers failed to detect a central tendency to deny knowledge, and rates of knowledge attribution reached nearly 60% in some conditions, even though the same participants tended to deny knowledge for other cases (Wright 2010).

The primary purpose of those first two studies of "fake barn" cases, it should be noted, was to look for order effects on knowledge attributions, so they used a within-subjects design, and they did not include control comparisons for the "fake barn" case specifically. No order effects were detected for the "fake barn" case (Swain, Alexander & Weinberg: 146; Wright 2010: 494, reporting $p = .086$ for the relevant comparison, which, contrary to what is suggested by the summary of results in the study's discussion section, is not statistically significant by conventional standards of interpretation). Several subsequent studies were designed to investigate knowledge judgments in "fake barn" cases specifically, and these studies included closely matched controls. The results have consistently shown that people tend to attribute knowledge to agents in such cases.

For example, one research team tested a case involving an agent, Gerald, who sees many things, including a cow and a real house amidst a large number of "house façades" (Colaćo, Buckwalter, Stich & Machery 2014). Participants then rated their agreement with a knowledge attribution on a 7-point Likert scale. In the control condition, participants rated whether Gerald knows that "he saw a cow." (The story did not mention anything about "fake cows.") In the experimental condition, participants rated whether Gerald knows that "he saw a house." Mean

score was higher in the control condition than in the experimental condition, indicating that the presence of façades depressed knowledge attribution. Nevertheless, mean response in the experimental condition was significantly above the scale's midpoint, indicating that people tended to attribute knowledge.

In another series of studies, people tended to attribute knowledge that an albino "vervet monkey" was in a tree even when it "was surrounded by" visibly indistinguishable "snow monkeys" (Turri 2016c). Again results from a closely matched control condition, in which the vervet monkey was a different color, suggesting that nearby, visibly similar distractors depress knowledge attribution. Nevertheless, even when there were nearby "fakes," attributing knowledge remained the central tendency.

Researchers also recently tested this pair of cases:

> (Control/Experimental) 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. 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. Despite all the [cheap barns/fakes] around, Sarah is in fact looking at the one [expensive/real] barn in the area. 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?" (Turri 2016b: 762)

Participants selected from one of four options that "best describes Sarah":

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.

(All options were randomly rotated and participants never saw numerical labels.) The overwhelming majority selected "She knows that it's a barn, and she should tell her son that it's a barn," in both the control (94%) and experimental (83%) conditions, with no statistically significant difference between them. These same basic finding was replicated using a 7-point Likert scale to collect knowledge judgments on their own (i.e. without being paired with assertability ratings in a series of conjunctions) (Turri 2017b).

Participants continued to attribute knowledge even when Sarah was explicitly described as incorrectly classifying four structures as barns before correctly classifying a fifth structure (Turri 2017b). That is, even in a "multiple-iteration fake barn case" where the agent first repeatedly encounters and misclassifies the first several "fakes" she sees, people tend to attribute knowledge when the agent correctly classifies the real barn she sees.

Interestingly, in a closely matched control condition that did not mention the existence of fakes or any other factor that would cause the agent to misidentify a non-barn as a barn, an importantly different pattern was observed: people tended to deny knowledge when the agent correctly classified the real barn on the fifth try. In other words, the presence of nearby fakes could actually prevent iterated errors from depressing knowledge judgments. Regression and causal-search analyses suggested that this effect was due to participants making different inferences about whether the agent was able to detect barns in the two conditions. When the errors were

plausibly due to the fakes, people still thought the agent was able to detect barns, and when she got the correct answer through that ability, they attributed knowledge (compare Turri 2016d). By contrast, when no information presented could plausibly explain the errors, participants inferred that the agent lacked the relevant perceptual ability, so they did not attribute knowledge.

At this point, it is clear that there is a central tendency to attribute knowledge in "fake barn" cases. Thus a founding assumption of this episode in professional epistemologists' fascination with "Gettier" cases — namely, that there such a tendency — is false. One might conjecture that this is an instance where "trained professionals" tend to have very different intuitions than untrained laypeople. But the evidence does not support that either: when researchers tested epistemologists' judgments about cases with a "fake barn" structure, most attributed knowledge (Horvath & Wiegmann 2016).

Conclusion

Despite a potential bright spot here or there, the theoretical literature on Gettier's cases, "Gettier" cases, and "the Gettier problem" is an embarrassing, confusing, and unproductive mess, for several different reasons. Many of the cases proposed in the literature, including Gettier's originals, are unsuitable to provide useful information because of their poor construction and prejudiced presentation. The theory that the cases supposedly undermine, the allegedly "traditional" theory of knowledge as "justified true belief," was not the traditional theory and, in fact, might not have been held by any influential philosopher before the mid-twentieth century. Moreover, prior to

Gettier, many philosophers had proposed simpler, clearer, and therefore better cases apt to teach the same lesson that Gettier's cases allegedly taught, so the intense focus on Gettier's later, inferior cases is both counterproductive and unfair. Additionally, philosophers have made unsupported empirical claims about how people tend to judge "Gettier" cases, and they have relied on these claims to motivate and evaluate theories, but empirical research has shown that some of those empirical claims are false. Philosophers are also apparently misinformed about how philosophers tend to judge some "Gettier" cases, with recent studies finding that most epistemologists attribute knowledge in cases that were allegedly intuitive cases of ignorance. Relatedly, perhaps, philosophers have incompetently curated the genre of thought experiments variously counted as "Gettier" cases: experimental studies have shown that knowledge attribution in different "Gettier" cases varies from lower than 20% to higher than 80%.

If judgments about "Gettier" cases vary this widely — from patterns resembling paradigmatic ignorance to patterns resembling paradigmatic knowledge — then "Gettier case" is a theoretically useless category. The fact that something is a "Gettier" case is consistent with its being both overwhelmingly judged knowledge and overwhelmingly judged ignorance. Such findings demonstrate the importance of including control conditions and tight manipulations using closely matched stimuli. Doing these things allows us to make meaningful comparisons and responsibly interpret results, which in turn sanctions (defeasible) inferences about features of people's knowledge concept. It can also prevent us from grouping wildly heterogeneous cases into theoretically useless categories.

So rather than dwell at length on "the Gettier problem" and try to "radically" adjust our

views about what the "distinctive" feature of Gettier's cases was (Hetherington 2016), and rather than try to "vindicate the tradition [sic] by showing JTB to be *almost* right" (Turri 2012b: 257), a much better way forward would be for epistemologists, individually and collectively, to stop crediting Gettier for an idea that was not originally his, stop repeating lies about the historical importance of the theory he criticized, stop ignoring the irredeemable faults of his original cases, and be much more discerning and humble about curating genres of thought experiment, identifying central tendencies in judgments about particular cases, and drawing theoretical conclusions based on such (alleged) tendencies. In short, end the malpractice.

Additionally, in light of the fact that mainstream philosophical research in this area has been guilty of grave shortcomings for so long, we should not take seriously vague, self-flattering worries about whether ordinary people can "competently assess" various cases, where the supposed standard of competence is the alleged central tendency in philosophers' judgments about the matter (Turri 2013). Philosophers should be more careful before proclaiming consensus on an issue and more mindful of the possibility that they lack unproblematic transparent access to their intuitions and attitudes about cases, let alone those of others. Relatedly, potentially fruitful lines of research in the psychology and sociology of philosophy could explore the extent to which an appearance of disciplinary "consensus" results from mechanisms of, at the very least, dubious intellectual legitimacy, such as thought-experimenter bias, gatekeeper effects, selection effects, false consensus effects, conformism, and motivated reasoning (for references and further discussion, see Turri 2016a).

Accordingly, I propose that after decades of generating much heat and little if any light, we

should close the book on this long and sad chapter of contemporary anglophone theoretical epistemology. In the aftermath, if for some unexpected reason, revisiting Gettier's cases or "the Gettier problem" merits our attention, then, by all means, let us do so, but this time much more seriously and honestly.

By contrast, in just a few short years, experimental work on "Gettier" cases has generated a range of informative, replicated results, which have begun answering questions about the ordinary knowledge concept. It is a good bet that further progress lies in this direction, so this sort of work should be welcomed and encouraged. The fact that previous results pertained to "Gettier" cases is incidental, however, because that is not a meaningful category. We would be better off simply dropping the eponymous labeling and instead focusing on the underlying factors that actually affect knowledge judgments, such as perceptions of ability, luck, deception, or interference, to name just a few.

Some philosophers might be tempted to insist that empirical results are irrelevant to their research project, which is "pursuing only the minutiae of a concept possessed" by "English-speaking philosophers" but not by other people (Lycan 2006: 165). Philosophers should resist this temptation because it is an empirical question which concepts philosophers possess. To the extent that they do succumb, philosophers should not be surprised when people stop ignoring their navel-gazing only to malign it.

Of course, studying a knowledge concept is not the only project one might be interested. Instead, one might wish to study knowledge itself, a real cognitive relation that often obtains between minds and facts, which our concepts could mischaracterize. In this chapter, I have not

concerned myself with this other project because, based on the methodology they utilize, it is perfectly clear that neither Gettier's paper nor most philosophical writings addressing "the Gettier problem" provide serious evidence about the nature of knowledge. They study thought experiments and argue about implications of stipulated verdicts for artificially precise definitions of knowledge; they do not study minds, mental states, or cognitive processes that actually generate, store, transmit, or utilize knowledge. It is equally clear, it should be added, and as would be happily acknowledged by researchers working in the field, that the experimental research reviewed here is informative regarding only a small fraction of what occurs in people's minds, namely, that pertaining to their knowledge concept and the processing of knowledge judgments in relation to a limited set of factors.

Acknowledgments — For helpful feedback and discussion, I thank Stephen Hetherington and Angelo Turri. This research was supported by the Canada Research Chairs Program and the Social Sciences and Humanities Research Council of Canada.

References

Beebe, J. R., & Buckwalter, W. (2010). The epistemic side-effect effect. Mind & Language, 25(4), 1–25.

Beebe, J. R., & Shea, J. (2013). Gettierized Knobe effects. Episteme, 10(03), 219–240. http://doi.org/10.1017/epi.2013.23

Blouw, P., Buckwalter, W., & Turri, J. (2017). 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.

Buckwalter, W. (2013). Gettier made ESEE. Philosophical Psychology, 1–16. http://doi.org/10.1080/09515089.2012.730965

Chisholm, R. M. (1989). Theory of Knowledge (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.

Church, I. M. (2013). Manifest failure failure: the Gettier problem revived. Philosophia, 41(1), 171–177. http://doi.org/10.1007/s11406-013-9418-5

Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. Episteme, 11(02), 199–212. http://doi.org/10.1017/epi.2014.7

Cullen, S. (2010). Survey-Driven Romanticism. Review of Philosophy and Psychology, 1(2), 275–296. http://doi.org/10.1007/s13164-009-0016-1

Dancy, J. (1985). Introduction to contemporary epistemology. Malden, Mass.: Blackwell.

Dutant, J. (2015). The legend of the justified true belief analysis. Philosophical Perspectives, 29, 95–145.

Feldman, R. (2003). Epistemology. Upper Saddle River, NJ: Prentice Hall.

Friedman, O., & Turri, J. (2015). Is probabilistic evidence a source of knowledge? Cognitive Science, 39(5), 1062–1080. http://doi.org/10.1111/cogs.12182

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

Goldman, A. I. (1967). A causal theory of knowing. Journal of Philosophy, 64(12), 357–372.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. Journal of Philosophy, 73(20), 771–791.

Greco, J. (2010). Achieving knowledge: A virtue-theoretic account of epistemic normativity. Cambridge: Cambridge University Press.

Hetherington, S. (2016). Knowledge and the Gettier problem. Cambridge: Cambridge University Press.

Horvath, J., & Wiegmann, A. (2016). Intuitive expertise and intuitions about knowledge. Philosophical Studies, 1–26. http://doi.org/10.1007/s11098-016-0627-1

Jackson, F. (1998). From metaphysics to ethics. Oxford: Oxford University Press.

Jackson, F. (2011). On Gettier Holdouts. Mind & Language, 26(4), 468–481.

Jacquette, D. (1996). Is nondefectively justified true belief knowledge? Ratio, 9(2), 115–127.

Kaplan, M. (1985). It's not what you know that counts. Journal of Philosophy, 82(7), 350–363.

Kim, M., & Yuan, Y. (2015). No cross-cultural differences in Gettier car case intuition: a replication study of Weinberg et al. 2001. Episteme, 12(03), 355–361. http://doi.org/10.1017/epi.2015.17

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

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.

Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., et al. (2015). Gettier across cultures. Nous.

Machery, E., Stich, S., Rose, D., Alai, M., Angelucci, A., Berniūnas, R., et al. (2017). The Gettier

intuition from South America to Asia. Journal of the Indian Council of Philosophical Research, 34(3), 517–541..

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. http://doi.org/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.

Plantinga, A. (1993). Warrant: the current debate. Oxford: Oxford University Press.

Powell, D., Horne, Z., Pinillos, A. & Holyoak, K. (2013). Justified true belief triggers false recall of knowing. Proceedings of the 35th Annual Conference of the Cognitive Science Society. Austin, TX.

Pritchard, D. (2005). Epistemic luck. New York : Oxford University Press.

Seyedsayamdost, H. (2015). On normativity and epistemic intuitions: failure of replication. Episteme, 12(01), 95–116. http://doi.org/10.1017/epi.2014.27

Sosa, E. (1991). Knowledge in Perspective. Cambridge: Cambridge University Press.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. Cognition, 124(3), 272–283. http://doi.org/10.1016/j.cognition.2012.05.017

Starmans, C., & Friedman, O. (2013). Taking “know” for an answer: A reply to Nagel, San Juan, and Mar. Cognition, 129(3), 662–665. http://doi.org/10.1016/j.cognition.2013.05.009

Swain, S., Alexander, J., & Weinberg, J. M. (2008). The Instability of Philosophical Intuitions:

Running Hot and Cold on Truetemp. Philosophy and Phenomenological Research, 84(1), 138–155.

Turri, J. (2012a). In Gettier's wake. In S. Hetherington (Ed.), Epistemology: The key thinkers (pp. 214–229). London: Continuum.

Turri, J. (2012). Is knowledge justified true belief? Synthese, 184(3), 247–259. http://doi.org/10.1007/s11229-010-9773-8

Turri, J. (2013). A conspicuous art: putting Gettier to the test. Philosophers' Imprint, 13(10), 1–16.

Turri, J. (2014). The problem of ESEE knowledge. Ergo, 1(4), 101–127.

Turri, J. (2015a). An open and shut case: epistemic closure in the manifest image. Philosophers' Imprint, 15(2), 1–18.

Turri, J. (2015b). Skeptical appeal: the source-content bias. Cognitive Science, 39(2), 307–324. http://doi.org/10.1111/cogs.12153

Turri, J. (2016a). Knowledge judgments in “Gettier” cases. In J. Sytsma & W. Buckwalter (Eds.), A companion to experimental philosophy (pp. 337–348). Wiley-Blackwell.

Turri, J. (2016b). Knowledge and assertion in “Gettier” cases. Philosophical Psychology, 29(5), 759–775. http://doi.org/10.1080/09515089.2016.1154140

Turri, J. (2016c). Vision, knowledge, and assertion. Consciousness and Cognition, 41(C), 41–49. http://doi.org/10.1016/j.concog.2016.01.004

Turri, J. (2016d). A new paradigm for epistemology: from reliabilism to abilism. Ergo, 3(8), 189–231.

Turri, J. (2017a). Experimental, cross-cultural, and classical Indian epistemology. Journal of Indian Council of Philosophical Research, 1–16. http://doi.org/10.1007/s40961-017-0117-7

Turri, J. (2017b). Knowledge attributions in iterated fake barn cases. Analysis, 77(1), 104–115. http://doi.org/10.1093/analys/anx036

Turri, J. (unpublished ms). A peculiar and perpetual tendency: an asymmetry in knowledge attribution for affirmations and negations. University of Waterloo.

Turri, J. (in press). Knowledge attributions and lottery cases: a review and new evidence.

Turri, J., & Buckwalter, W. (2017). Descartes"s schism, Locke"s reunion: completing the pragmatic turn in epistemology. American Philosophical Quarterly, 54(1), 25–46.

Turri, J., & Friedman, O. (2014). Winners and losers in the folk epistemology of lotteries. In J. R. Beebe (Ed.), Advances in experimental epistemology (pp. 45–69). London: Bloomsbury.

Turri, J., Buckwalter, W., & Blouw, P. (2014). Knowledge and luck. Proceedings of the 36th Annual Conference of the Cognitive Science Society, 1958–1963.

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

Turri, J., Buckwalter, W., & Rose, D. (2016). Actionability judgments cause knowledge judgments. Thought: a Journal of Philosophy, 5(3), 212–222. http://doi.org/10.1002/tht3.213

Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. Philosophical Topics, 29(1&2), 429–460.

Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. Cognition, 115(3), 491–503. http://doi.org/10.1016/j.cognition.2010.02.003

Zagzebski, L. T. (1996). Virtues of the mind: an inquiry into the nature of virtue and the ethical foundations of knowledge. Cambridge: Cambridge University Press.