

Knowledge and luck

John Turri · Wesley Buckwalter · Peter Blouw

Published online: 9 July 2014
© Psychonomic Society, Inc. 2014

Abstract Nearly all success is due to some mix of ability and luck. But some successes we attribute to the agent's ability, whereas others we attribute to luck. To better understand the criteria distinguishing credit from luck, we conducted a series of four studies on knowledge attributions. Knowledge is an achievement that involves reaching the truth. But many factors affecting the truth are beyond our control, and reaching the truth is often partly due to luck. Which sorts of luck are compatible with knowledge? We found that knowledge attributions are highly sensitive to lucky events that change the explanation for why a belief is true. By contrast, knowledge attributions are surprisingly insensitive to lucky events that threaten, but ultimately fail to change the explanation for why a belief is true. These results shed light on our concept of knowledge, help explain apparent inconsistencies in prior work on knowledge attributions, and constitute progress toward a general understanding of the relation between success and luck.

Keywords Knowledge attribution · Luck · Social cognition

Reasoning about what people know is central to our lives and is often essential for predicting and evaluating human actions. Whether someone knows that a certain action is impermissible affects whether we excuse or punish them for it (Duff, 1990;

Electronic supplementary material The online version of this article (doi:10.3758/s13423-014-0683-5) contains supplementary material, which is available to authorized users.

J. Turri (✉)
Department of Philosophy, University of Waterloo,
Waterloo, ON, Canada
e-mail: john.turri@gmail.com

W. Buckwalter · P. Blouw
University of Waterloo, Waterloo, ON, Canada

Hart, 1959; Malle & Nelson, 2003). Moreover, if you know something, then you are entitled to base actions on it and communicate it to others. By contrast, if you do not know something, then it seems that you should be more cautious about basing actions on it or communicating it to others (Turri & Buckwalter, 2014; Fantl & McGrath, 2009; Hawthorne, 2004; Stanley, 2005; Turri, 2011a, 2014b; Unger, 1975; Williamson, 2000). Indeed, researchers have recently argued that knowledge sets the standard for the two main forms of human pedagogy—assertion (telling someone a fact) and instructional demonstration (showing someone how something is done)—which makes knowledge an essential ingredient of human culture and civilization (Buckwalter & Turri, *in press*).

Knowledge is an achievement that involves reaching the truth through cognitive ability (Aristotle 1941; Greco, 2010; Reid, 1785/2002; Sosa, 2007). Nearly all human achievement, including knowledge, is due to some mix of ability and luck. Ludwig Wittgenstein had in mind this inevitable role of luck when he wrote, “It is always by favour of Nature that one knows something” (1975, §505). But not just any sort of luck is compatible with knowing (Plato, 1992; Russell, 1948, p. 113). For example, blindly guessing the truth does not yield knowledge. But beyond this obvious starting point, which sorts of luck are viewed as being compatible with knowledge?

Insightful work in philosophy and psychology has focused on one or another aspect of the relationship between knowledge and luck (Beebe & Shea, 2013; Buckwalter, 2013; Cullen, 2010; Engel, 1992; Gettier, 1963; Goldman, 1976; Harman, 1973; Nagel, San Juan, & Mar, 2013; Pritchard, 2005; Starmans & Friedman, 2012, 2013; Unger, 1968; Weinberg, Nichols, & Stich, 2001; Wright, 2010), but the relationship has not been systematically investigated from a psychological perspective. Other psychological work has investigated how a neighboring concept related to luck, “deviant causation” (Searle, 1983), affects people's attributions of

intentional action and judgments of moral responsibility (Knobe, 2003; Malle, 2006; Malle & Knobe, 1997; Mele & Cushman, 2007; Pizarro, Uhlmann, & Bloom, 2003). Recent work has also identified at least one form of luck, broadly related to deviant causation, that does not seem to diminish knowledge attribution (Starmans & Friedman, 2012), although subsequent work has claimed to undermine these findings (Nagel et al., 2013).

Philosophers and psychologists have developed many thought experiments intended to examine intuitions about the relationship between knowledge and luck (e.g., Beebe & Buckwalter, 2010; Shope, 1983; Starmans & Friedman, 2012; Turri, 2014a; Wright, 2010). These thought experiments have varied widely and have focused on different features that different authors have considered relevant to luck. There are almost as many thought experiments as there are authors who have written on the topic. The variety of cases exhibits three important structural differences. 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. 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. At 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 closely the originally perceived truth-maker and backup truth-maker resemble one another. Sometimes they very closely resemble one another, but at other times they differ greatly.

Consider four specific cases that have loomed large in recent debates.

(BARN) Henry and his son are driving through the country. Henry pulls over to stretch his legs, and while doing so tells his son a list of items currently in view along the roadside. “That’s a tractor. That’s a combine. That’s a horse. That’s a silo. And that’s a fine barn,” Henry added, pointing to the nearby roadside barn, which is indeed very fine. But Henry is unaware that the locals recently replaced nearly every barn in the county with papier-mâché fake barns. Henry is looking at the one real barn in the whole county. (Adapted from Goldman, 1976).

(PEN) Katie is in her locked apartment writing a letter. She puts the letter and her blue Bic pen down on her coffee table. Then she goes into the bathroom to take a shower. As Katie’s shower begins, two burglars silently break into the apartment. One burglar takes Katie’s blue Bic pen from the table. But the other burglar absent-mindedly leaves his own identical blue Bic pen on the coffee table. Then the burglars leave. Katie is still in the

shower and did not hear anything. (Starmans & Friedman, 2012, p. 276).

(HUSBAND) Mary enters the house and looks into the living room. A familiar appearance greets her from her husband’s chair. She thinks, “My husband is sitting in the living room,” and then walks into the den. But Mary misidentified the man in the chair. It’s not her husband but his brother, whom she had no reason to think was even in the country. However, her husband is seated along the opposite wall of the living room, out of Mary’s sight, dozing in a different chair. (Turri, 2011b, p. 2; adapted from Zagzebski, 1996; cf. Chisholm, 1989).

(BARCELONA) Smith has strong evidence that Jones owns a Ford. Smith has another friend, Brown, of whose whereabouts he is totally ignorant. On the basis of his evidence about Jones, Smith accepts the proposition that “Either Jones owns a Ford, or Brown is in Barcelona,” even though he has no idea where Brown is. It turns out that Jones does not own a Ford and is presently driving a rented car. However, by the sheerest coincidence and entirely unknown to Smith, Brown is traveling in Barcelona. (Adapted from Gettier, 1963, pp. 122–123).

These cases exemplify all of the important differences noted above. First, in the BARN and PEN stories, the agent initially perceives a truth-maker, but this does not happen in either HUSBAND or BARCELONA. Second, in BARN the agent’s perceptual relation to the truth-maker remains intact throughout, despite the threat, made salient to the reader, that he could have perceptually misidentified the relevant item. By contrast, in PEN the truth-maker is not only threatened but also disrupted and replaced, thereby changing the explanation for why the agent’s belief is true. Third, in HUSBAND, even though the agent does not perceive the truth-maker, her belief is nevertheless made true by a state of affairs that very closely resembles the state of affairs that she thinks makes her belief true. By contrast, in BARCELONA the agent’s belief is made true by a state of affairs that differs greatly from the one that he thinks makes his belief true. In other words, the “backup” truth-maker in HUSBAND is very similar to what the agent had in mind, whereas in BARCELONA it is very dissimilar.

Theorists have argued at length over the intuitively correct verdict in each of these cases—does the agent *know* or *only believe* the truth?—and over the best explanation for that verdict (for an overview, see Shope, 2002). We will not enter into the theoretical debate here, at least not directly (see the General Discussion for ways that our results could indirectly shed light on the theoretical debate). Instead, we are interested in whether people’s ordinary judgments about knowledge are sensitive to the three luck-related factors identified in the previous paragraph: *threat*, *disruption*, and *replacement*. More specifically, we tested the effect of the following things on ordinary knowledge attributions: (i) A threat to the detection

of truth: When an agent perceptually detects a truth-maker, what is the effect of a salient threat to her ability to detect it? (ii) The disruption of a truth-maker: What is the effect of a successful threat that fundamentally changes the explanation for why her belief is true? (iii) The resemblance between the initial truth-maker and the backup truth-maker that replaces it: What is the effect of making them very similar or dissimilar?

The results of our investigations will highlight some general lessons about the ordinary concept of knowledge and provide a framework for integrating prior psychological findings on knowledge attributions and resolving some potentially troubling inconsistencies in the literature. Ultimately, it turns out that there is no true and general lesson to be drawn about the relationship between knowledge and luck. Instead, there are several more specific lessons about the relationship between knowledge and the different luck-related factors, each of which affects knowledge attributions in interesting and importantly different ways. Our results also shed light on the potential theoretical usefulness of a peculiar class of cases, known as “Gettier cases,” that have recently generated controversy in the psychological literature (Nagel, Mar, & San Juan, 2013a, b; Starmans & Friedman, 2012, 2013; Turri, 2013). In the General Discussion, we will examine in greater detail the relationship between our findings and prior work on these issues.

Experiment 1

This first experiment tested the effect on knowledge attributions of a salient but failed threat to the truth of a perceptual judgment. In many such cases, the agent is clearly lucky that the threat fails to prevent the formation of a true belief. Some philosophers and psychologists have claimed that such a threat obviously suffices to prevent one from gaining knowledge (Goldman, 1976; Nagel et al., 2013b; Pritchard, 2005), but this verdict has been disputed (Colaço, Buckwalter, Stich, & Machery, 2014; Lycan, 2006; Turri, 2011b). To test which side of the dispute better captures the ordinary view of knowledge, we compared rates of knowledge attributions in response to three different stories in a between-subjects experiment. In the first story, an agent forms a true belief based on perceiving the truth-maker, and nothing threatens to disrupt the truth-maker. In the second story, an agent forms a true belief based on perceiving the truth-maker, and there is a salient but failed threat to his ability to detect it. In the third story, the threat succeeds and prevents the agent from forming a true belief: there is no truth-maker for the agent’s belief. We expected that the first story would be viewed as a paradigm case of knowledge and that the third story will be viewed as a paradigm case of ignorance, in which case they will serve as useful control comparisons in relation to the second story. Our main question was whether the rate

of knowledge attribution in response to the second story would more closely resemble the rate in the first story (knowledge control) or in the third story (ignorance control). We expected that it would closely resemble the rate for the first story and be significantly higher than the rate for the third story.

Method

Participants A group of 135 participants (18–59 years of age, mean age = 29.1 years; 94 male, 41 female) were recruited and tested using an online platform (Amazon Mechanical Turk [AMT] + Qualtrics) and compensated \$0.30 for approximately 2 min of their time. Participation was restricted to United States residents, and 94 % reported English as a native language. Participants were recruited and compensated similarly for all subsequent experiments. Repeat participation, within and across experiments, was prevented by allowing each AMT worker to complete a task only once and by screening the data for responses from participants with identical Worker IDs across experiments. We excluded data from 15 participants who failed comprehension questions. Including data from these participants did not affect the results reported below.

Materials and procedure Participants were randomly assigned to one of three conditions and read a single story in a between-subjects design. The three conditions were No Threat (knowledge control), Threat, and No Detection (ignorance control).

The basic storyline featured an ecologist, Darrel, collecting data on a (fictional) species, the red speckled ground squirrel. The story for each condition began in the same way.

Darrel is an ecologist collecting data on red speckled ground squirrels in Canyon Falls national park. The park is divided into ten zones and today Darrel is working in Zone 3. ¶¹ While scanning the river valley with his binoculars, Darrel sees a small, bushy-tailed creature with distinctive red markings on its chest and belly. The red speckled ground squirrel is the only native species with such markings. Darrel records in his journal, “At least one red speckled ground squirrel in Zone 3 today”.

The stories for the three conditions differed in their final paragraph. Here are the respective final paragraphs:

(No Threat) Ecologists are unaware that a complex network of aquifers recently began drying up in parts of the park. These aquifers carry vital nutrients to the

¹ Indicates paragraph break on participant’s screen.

trees and other forms of plant life that support the squirrels. And the aquifers in the river valley running through Zone 3 are no exception. The animal Darrel is looking at is indeed a thirsty red speckled ground squirrel.

(Threat / No Detection) Ecologists are unaware that a non-native species of prairie dog recently began invading the park. These prairie dogs also have red markings on their chest and belly. When these prairie dogs tried to invade Zone 3, the red speckled ground squirrels were unable to completely drive them away. [Still, / And] the animal Darrel is looking at is indeed [a red speckled ground squirrel / one of the prairie dogs].

Participants then answered three questions. Participants first responded to a knowledge question standardly used in the study of knowledge attribution (Friedman & Turri, *in press*; Nagel et al., 2013b; Starmans & Friedman, 2012; Turri, 2013).

Darrel _____ that there is at least one red speckled ground squirrel in Zone 3 today [knows / only believes].

Participants were then taken to a separate screen where they answered a comprehension question and a question about what it was reasonable for Darrel to believe.

Darrel is looking at a _____. [ground squirrel/prairie dog]
It is _____ for Darrel to think that he is looking at a red speckled ground squirrel [reasonable/unreasonable].

Participants were not allowed to go back to a previous page and change their answer. Questions were always asked in the same order, and response options were rotated randomly. After testing, participants filled out a brief demographic questionnaire. These same basic procedures were followed in all subsequent experiments reported in this article.

Results and discussion

Preliminary analysis revealed no effect of participant gender on responses to the knowledge question, either in the aggregate, Fisher's exact test $p = .262$, or for any of the three stories, Fisher's $ps \leq .394$, so the analyses that follow are collapsed across genders. The same is true for all subsequent experiments. This is consistent with prior work on knowledge attributions that had shown no effect of demographic variables (Nagel et al., 2013b; Turri, 2013; Wright, 2010).

As predicted, assignments to conditions affected the rates of knowledge attribution, $\chi^2(df = 2, N = 135) = 39.63, p < .001$, Cramer's $V = .542$ (all tests are two-tailed unless otherwise noted). (See Table 1.) Pairwise comparisons detected no difference in knowledge attributions between the No Threat and Threat conditions, Fisher's $p = .164$, n.s., and a large difference between Threat and No Detection, Fisher's

Table 1 Experiment 1: Percentages of participants attributing knowledge and agreeing that it was reasonable for the protagonist to think that he was looking at an object of the relevant sort

	No Threat	Threat	No Detection
Knows	81 %	67 %	16 %
Reasonable	98 %	94 %	87 %

$p < .001$, Cramer's $V = .509$. As expected, No Threat was overwhelmingly viewed as a case of knowledge, and No Detection was overwhelmingly viewed as a case of ignorance. Binomial tests revealed that knowledge attribution was above chance rates in the No Threat, $p < .001$, test proportion = .5, and in the Threat, $p = .021$, conditions, whereas it was significantly below chance in the No Detection condition, $p < .001$. Condition did not affect whether people said that it was reasonable for the protagonist to think he was looking at an object of the relevant sort, $\chi^2(df = 2, N = 135) = 4.49, p = .106$, n.s., so the differences in knowledge attribution can not be due to perceived differences in what it was reasonable for the protagonist to believe. (See Table 1.)

The results support the view that a salient but failed threat to the truth of a judgment does not significantly affect whether it is viewed as knowledge. When the threat *failed* to prevent Darrel from detecting the truth, participants attributed knowledge at rates exceeding chance, and indeed, at rates not significantly different from a case in which no such threat was mentioned. By contrast, when the threat *succeeded* in preventing Darrel from detecting the truth, participants overwhelmingly declined to attribute knowledge. The luck involved in such a threat failing is viewed as being fully consistent with knowledge.

Experiment 2

This second experiment tested the effect on knowledge attributions of an unnoticed change in the explanation for why the agent's belief is true. More specifically, it tested the effect of an unnoticed disruption of the initial truth-maker and its replacement with a backup. (It is important to note that disrupting the initial truth-maker does not automatically make the belief false, for it could turn out that some other state of affairs could serve as a backup truth-maker.) Many researchers have claimed that the luck involved with such disruption and replacement-by-backup prevents one from knowing (Lehrer, 1965; Sosa, 2007), but this verdict has been disputed (Hetherington, 1998, 1999, 2013; Sartwell, 1991, 1992; Starmans & Friedman, 2012). To test which side of the dispute better captures the ordinary view of knowledge, we compared rates of knowledge attributions in response to three different

stories in a between-subjects experiment. In the first story, an agent forms a true belief based on perceiving the truth-maker, and nothing threatens to disrupt the truth-maker. In the second story, the agent forms a true belief based on perceiving the truth-maker, and there is an unnoticed change in what makes the belief true: the initial truth-maker is disrupted and replaced by a backup. In the third story, the agent fails to detect the truth initially, and what actually makes the belief true goes unnoticed. Again we expected that the first story would be viewed as a paradigm case of knowledge and that the third story would be viewed as a case of ignorance. Our main question was whether the rate of knowledge attribution in response to the second story would more closely resemble the rate in the first story (knowledge control) or in the third story (ignorance control). We expected the rate for the second story to be lower than for the first and higher than for the third, but we also expected it to more closely resemble the rate for the first.

Method

Participants A group of 141 new participants (18–66 years of age, mean age = 31.2 years; 84 male, 57 female, 99 % reporting English as a native language) were tested. We excluded data from six participants who failed a comprehension question and from three participants whose Worker ID matched a participant's ID from the previous study. Including data from these participants did not affect the results reported below.

Materials and procedure Participants were randomly assigned to one of three conditions and read a single story in a between-subjects design. The three conditions were Normal Detection (knowledge control), Replacement, and No Detection (ignorance control).

The basic storyline again featured the ecologist Darrel collecting data on the red speckled ground squirrel. This time the story covered a longer time period, introduced a second character who asked Darrel a question, and involved Darrel relying on memory as well as perception. The story for each condition began the same way.

Darrel is an ecologist collecting data on red speckled ground squirrels in Canyon Falls national park. The park is divided into ten zones and today Darrel is working in Zone 3. ¶ While scanning the river valley with his binoculars, Darrel sees a bushy-tailed creature with distinctive red markings on its chest and belly. The red speckled ground squirrel is the only native species with such markings. Soon Darrel packs up his gear and hikes back to base camp. ¶ When Darrel returns to camp, his colleague says, “A reporter is going to do a story on local wildlife and she needs some video footage of a red

speckled ground squirrel. Are there any in Zone 3?” Darrel replies, “Yes, there is at least one red speckled ground squirrel in Zone 3.”

The stories for the three conditions differed in their final paragraph. In the Normal Detection condition, the animal that Darrel saw was a female red speckled ground squirrel, and while Darrel was hiking back to camp, a male squirrel migrated into Zone 3 and joined her. In the Replacement condition, the animal that Darrel saw was a female red speckled ground squirrel, but while Darrel was hiking back to camp, the female that he saw migrated out of Zone 3 and never returned. However, a different female migrated into Zone 3 and nested in the river valley, so there was a red speckled ground squirrel in Zone 3 after all. In the No Detection condition, the animal that Darrel saw was actually a member of an invasive species of prairie dog, so that he did not detect a red speckled ground squirrel in Zone 3 that day. However, as in the story for the Replacement condition, while Darrel was hiking back to camp, a female red speckled ground squirrel migrated into Zone 3 and nested in the river valley, so that a red speckled ground squirrel was in Zone 3 after all. The complete texts for all of the stories are included in the [supplemental materials](#).

Participants then answered three questions. Participants first responded to the same knowledge question as in Experiment 1. Participants were then taken to a separate screen where they answered a comprehension question and a question about what it was reasonable for Darrel to believe. These were similar to the final two questions in Experiment 1, except that this time they were in the past tense rather than the present.

Results and discussion

As predicted, assignments to conditions affected the rates of knowledge attribution, $\chi^2(df=2, N=141) = 40.16, p < .001$, Cramer's $V = .534$. (See Table 2.) Pairwise comparisons revealed that knowledge attribution was higher in Normal Detection than in Replacement, Fisher's $p = .017$, Cramer's $V = .254$, and higher in Replacement than in No Detection, Fisher's $p < .001$, Cramer's $V = .427$. As expected, Normal Detection was overwhelmingly viewed as a case of knowledge, and No Detection was overwhelmingly viewed as a case of ignorance. Binomial tests revealed that knowledge

Table 2 Experiment 2: Percentages of participants attributing knowledge and agreeing that it was reasonable for the protagonist to think that he saw an object of the relevant sort

	Normal Detection	Replacement	No Detection
Knows	88 %	66 %	23 %
Reasonable	98 %	100 %	98 %

attribution was above chance rates in Normal Detection, $p < .001$, test proportion = .5, and in Replacement, $p = .033$, whereas it was significantly below chance in No Detection, $p < .001$. Condition did not affect whether people said it was reasonable for the protagonist to think that he saw an object of the relevant sort, $\chi^2(df=2, N=141) = 1.124, p = .570$, n.s., so the differences in knowledge attribution again can not be due to perceived differences in what it was reasonable for the protagonist to believe. (See Table 2.)

The results support the view that an unnoticed replacement-by-backup does affect whether an belief is viewed as knowledge. When replacement does not occur, knowledge attribution is at ceiling, but when it does occur, knowledge attribution is significantly lower. Nevertheless, in this study such a replacement was widely viewed as being consistent with knowing: Most participants in the Replacement condition attributed knowledge, even though the belief was ultimately true because of the backup truth-maker. However, participants' willingness to attribute knowledge in the Replacement condition seemed to depend crucially on whether the initial truth-maker was detected, as demonstrated by the very low rates of knowledge attribution in the No Detection condition, for the difference between the Replacement and No Detection conditions lay in whether the protagonist detected an initial truth.

Experiment 3

Having seen that truth-maker replacement does affect knowledge attribution, in this experiment we tested the effect of similarity or dissimilarity of the initial and replacement truth-makers. We suspected that a highly dissimilar truth-maker could draw greater attention to the fact that replacement had occurred, which could in turn lead people to view the agent as being lucky and withhold the attribution of knowledge. To test this, we compared rates of knowledge attributions in response to three different stories in a between-subjects experiment. In both the first and second stories, the agent formed a true belief based on perceiving the truth-maker, and there was an unnoticed change in what made the belief true: The initial truth-maker was disrupted and replaced by a backup. In the first story, the backup truth-maker was very similar to the initial truth-maker, whereas in the second story it was dissimilar. In the third story, the agent failed to detect the truth initially and nothing made the belief true. The first story was almost identical to the story for the Replacement condition in Experiment 2, so we expected participants to tend to attribute knowledge in response to the first story in this experiment. We expected the third story to be viewed as a paradigm case of ignorance, and included it as a control. Our main question was whether the rate of knowledge attribution in response to the

second story would differ from that for the first story, and if so, whether it would more closely resemble the rate for the first or for the third story. We expected the rate for the second story to be lower than that for the first, but to be higher than that for the third.

Method

Participants A group of 576 new participants (18–71 years of age, mean age = 31 years; 385 male, 191 female, 95 % reporting English as a native language) were tested.² We excluded data from 18 participants who failed the comprehension question and six participants who had previously participated in a related study. Including data from these participants did not affect the results reported below.

Materials and procedure Participants were randomly assigned to one of three conditions and read a single story in a between-subjects design. The three conditions were Similar, Dissimilar, and No Detection (ignorance control). The basic storyline again featured the ecologist Darrel collecting data on the red speckled ground squirrel. The story for each condition began with the same common stem as the stories for Experiment 2, but the stories then differed in their final paragraph. In both the Similar and Dissimilar conditions, the animal that Darrel saw was a female red speckled ground squirrel, and while Darrell was hiking back to camp, the female that he saw had migrated out of Zone 3 and never returned. In the Similar condition, the backup truth-maker was very similar to the initial truth-maker: Another female squirrel migrated into Zone 3 and nested in the same river valley that Darrel had been in earlier. By contrast, in the Dissimilar condition, the backup truth-maker was noticeably dissimilar from the original: A male red speckled ground squirrel was hibernating unnoticed in a burrow all the way on the other side of Zone 3. In the No Detection condition, the animal that Darrel saw was actually a member of an invasive species of prairie dog, and there was no backup truth-maker: There simply had been no red speckled ground squirrel in Zone 3.

² The number of participants was higher in this experiment in response to suggestions made during the review process. A previous trial of this experiment produced results very similar to those reported below, including the nonreplication of the result from the Replacement condition in Experiment 2. This raised questions about replicability, so we ran the study with a much larger sample size. This trial was conducted using a predetermined sample size gleaned from a priori power analysis, rounding up to compensate for exclusions due to comprehension failures and potential repeat participation. A priori power analysis indicated that a sample size of 183 per group had sufficient power (.80) to detect a .15 difference in proportions (.65–.50) for independent groups on the dichotomous knowledge question. This was calculated using the software application G*Power, version 3.1.5 (Faul, Erdfelder, Lang, & Buchner, 2007). The calculation assumed the conventional alpha level of .05 and a beta level of .20.

The complete texts for all the stories are included in the [supplemental materials](#).

Participants then answered the same three questions as in Experiment 2: a knowledge question, a comprehension question, and a question about what it was reasonable for Darrel to believe. But since the answer to the reasonableness question did not matter in the first two experiments and response rates continued to be uniformly at or near ceiling (86 %–100 %), we omit further discussion of it for the sake of brevity.

Results and discussion

Assignments to conditions affected the rates of knowledge attribution, $\chi^2(df = 2, N = 576) = 94.97, p < .001$, Cramer's $V = .406$. (See Table 3) Pairwise comparisons revealed that knowledge attribution was, as predicted, higher in Similar than in Dissimilar conditions, Fisher's $p = .019$, Cramer's $V = .124$, and higher in Dissimilar than in No Detection conditions, Fisher's $p < .001$, Cramer's $V = .386$. As expected, No Detection was overwhelmingly viewed as a case of ignorance, with rates of knowledge attribution falling far below chance, binomial, $p < .001$, test proportion = .5. Knowledge attribution was also below chance in Dissimilar, binomial, $p = .031$. Knowledge attribution in the Similar condition did not differ from chance, binomial $p = .249$, and it was lower than the rate of knowledge attribution in the Replacement condition from Experiment 2, binomial $p = .002$, test proportion = .66. Given that the Replacement story from Experiment 2 was virtually identical to the Similar story in this experiment, we were surprised at this last difference. One possibility is that the population tends to attribute knowledge in such cases at rates slightly exceeding chance, and the divergent results that we observed were simply the result of random variation. Another possibility is that there are some important, currently unidentified individual differences in knowledge attributions, and that the sample from Experiment 2 included a greater proportion of people with a trait making them more tolerant of the relevant form of luck. (See Feltz & Cokely, 2012, p. 231, for a review of some evidence that personality traits predict judgments about particular cases used to support theories in moral psychology and philosophy.)

The results support the view that when disruption and replacement occurs, knowledge attribution can be affected by how similar or dissimilar the replacement truth-maker is to the original. People were more likely to attribute knowledge when the replacement was highly similar to the original than

when it was dissimilar. Cases involving dissimilar replacements are still viewed as being importantly different from paradigmatic cases of ignorance.

Experiment 4

Our final experiment was designed to replicate the findings from the first three experiments in a single study that used a very different cover story. The results not only replicated earlier findings, but also demonstrated that the findings generalize to other contexts and are not due to the use of a particular type of story. All of our earlier experiments had featured stories about nonhuman animals in natural outdoor settings. By contrast, the present experiment featured stories set in indoor social contexts involving owned artifacts and other human agents with malicious intentions.

Method

Participants A total of 813 new participants were tested (18–75 years of age, mean age = 31.25; 487 male, 326 female, 96 % reporting English as a native language). We excluded data from 85 participants who failed comprehension questions. Including data from these participants did not affect the results reported below.

Materials and procedure Participants were randomly assigned to one of seven conditions and read a single story. (See Table 4.) Each condition featured a different story. The basic storyline featured an agent, Emma, admiring jewelry in a fancy department store (based on a story in Nagel et al., 2013b). Emma purchased a stone from the diamond display, put it in her pocket, browsed for another minute, then left the store. The different versions of the story varied whether the stone was a real or a fake diamond, whether there was a threat to the stone remaining in Emma's pocket, whether the threat failed or succeeded, and whether any other stone also ended up in Emma's pocket. To adopt the terminology used in previous experiments, the different stories manipulated whether Emma detected an initial truth-maker, whether Emma's truth detection was saliently threatened, whether the threat succeeded in disrupting the initial truth-maker, and whether the backup truth-maker was highly similar or dissimilar to the initial.

In all of the stories, Emma purchased a stone from a jewelry store, put it in her pocket, and soon walked out of the store. In all of the stories that involved detection, the stone she purchased was a diamond. In all the stories that did not involve detection, the stone was a fake. In all of the stories that involved similar backup truth-makers, the backup truth-maker was that, one way or another, a diamond was put into

Table 3 Experiment 3: Percentages of participants attributing knowledge

	Similar	Dissimilar	No Detection
Knows	54 %	42 %	8 %

Table 4 Experiment 4: Description of stories used in the seven conditions of Experiment 4

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.

Emma's pocket before she left the store. In all of the stories that involved dissimilar backup truth-makers, the backup truth-maker was that a real diamond had secretly been sewn into Emma's pocket by a previous owner long ago. Table 4 describes the stories used for each of the conditions. The complete text of all of the stories is included in the [supplemental materials](#).

We note one residual issue from an earlier study. We had observed no statistically significant difference in knowledge attributions between the No Threat (knowledge control) condition and the Threat condition in Experiment 1, but the absolute difference was perhaps large enough (81 %/67 %) to warrant suspicion that a larger sampling would have revealed a difference. So we increased the number of participants per condition in order to gain statistical power. We also tested multiple representative versions of the stories for the Failed Threat and Detect Similar conditions as a precaution; we expected no difference between representative versions of the same type.

After reading the story, participants answered a series of comprehension questions about whether Emma had a diamond in her pocket when leaving the store, whether it was reasonable for Emma to think that there was a diamond in her pocket, and why Emma thought that there was a diamond in her pocket. All participants then answered the same test question about whether, as Emma left the store, she knew that a diamond was in her pocket. Questions were always asked in the same order, and response options were rotated randomly.

Four noteworthy differences in procedure distinguished this experiment from the first three. First, whereas in Experiments 1–3 we had asked the knowledge question first and then asked comprehension questions, in this experiment we asked the comprehension questions first and then the

knowledge question. Second, we included more comprehension questions this time, including one that required participants to explicitly identify Emma's reason for thinking that she had a diamond in her pocket. Third, we treated the reasonableness question as a comprehension question this time. Fourth, we also added explicit time-indexing to the questions, to minimize the possibility that participants would selectively interpret the questions as applying to different times in the story.

Results and discussion

First, we make three preliminary notes about the analyses that follow: (i) Knowledge attribution did not differ across the three versions of the Failed Threat story (80 %–88 %), $\chi^2(df = 2, N = 180) = 1.65, p = .438, n.s.$, so we collapsed these into a single group; (ii) knowledge attribution did not differ between the two versions of the story representative of category 3 (54 %/60 %), Fisher's $p = .497, n.s.$, so we collapsed these into a single group; (iii) gender did not affect the attribution rates (male, 54 %; female, 53 %), Fisher's $p = .886, n.s.$, so we collapsed across gender.

Assignments to conditions affected rates of knowledge attribution: 12 %–90 %, $\chi^2(df = 6, N = 813) = 242.21, p < .001$, Cramer's $V = 0.546$ (Fig. 1), and overall the results replicated the main findings from our first three experiments. First, the rates of knowledge attribution did not differ between Knowledge Control (90 %) and Failed Threat (83 %), Fisher's $p = .113, n.s.$, and far exceeded chance rates in both conditions, binomial $ps < .001$. This replicated a main finding from Experiment 1: A failed threat to truth detection does not significantly depress knowledge attribution and is viewed as being fully consistent with knowledge. This result is also in




Fig. 1 Experiment 4: Percentages of participants attributing knowledge across conditions. Except where nonsignificance is indicated, the significance for all comparisons is at the $p < .01$ level.

line with findings on related cases by Colaço et al. (2014). Second, knowledge attribution in Detect Similar (55 %) was significantly lower than in Knowledge Control, Fisher's $p < .001$, Cramer's $V = .367$, was nonsignificantly above chance, binomial $p = .195$, and was significantly higher than in Ignorance Control, Fisher's $p < .001$, Cramer's $V = .416$. This replicated the main findings from Experiment 2: An unnoticed truth-maker replacement did depress knowledge attribution, was widely, though not uncontroversially, viewed as consistent with knowledge, and was certainly viewed very differently from a clear case of ignorance. Third, rates of knowledge attribution differed significantly, depending on whether the replacement truth-maker was similar to or dissimilar from the initial truth-maker. This was true regardless of whether the agent had initially detected a truth-maker: Rates of knowledge attribution were higher in Detect Similar than in Detect Dissimilar (39 %), Fisher's $p = .016$, Cramer's $V = .156$, and they were higher in No Detect Similar (38 %) than in No Detect Dissimilar (19 %), Fisher's $p = .006$, Cramer's $V = .207$. This not only replicated a main finding from Experiment 3, but it also generalized that finding, for the finding in Experiment 3 was limited to a comparison involving cases of detection, whereas here we found that the difference extended to cases not involving detection.

Three other aspects of the results are worth noting explicitly. First, knowledge attribution exceeded chance rates only when the protagonist detected the truth and the truth-making relation was not disrupted (Knowledge Control and Failed Threat). Second, knowledge attribution fell below chance rates whenever the protagonist failed to detect the truth (No Detect Similar, No Detect Dissimilar, and Ignorance Control [12 %]). Third, a protagonist with a false belief was overwhelmingly viewed as being ignorant, even when the false belief was eminently reasonable (Ignorance Control). This

supports a widespread assumption in both theoretical epistemology (Buckwalter, *in press*; Hazlett, 2010; Turri, 2011c) and developmental research on the acquisition of cognitive vocabulary (Booth, Hall, Robison, & Kim, 1997; Shatz, Wellman, & Silber, 1983, pp. 318–319; Sodian 1988)—namely, that the concept of knowledge is “factive,” meaning that its application presupposes that the claim known is true. We think an interpretation based on factivity is the simplest and most consistent with prior findings on knowledge attributions. Nevertheless, some caution is warranted here, because the results are consistent with more complicated alternative interpretations of the data that reject factivity. We welcome and encourage work that explores such alternatives.

General discussion

Like any human achievement, knowledge is usually due to a mix of ability and luck. But the effect of luck on knowledge attributions is not well understood. We tested the effect of three luck-related factors on knowledge attribution—threat, disruption, and replacement—and found that the relation is complex and often surprising. We suggest that our investigation is profitably viewed as a case study in how people judge the relationship between success and luck. We focused on a centrally important class of cognitive evaluations, knowledge judgments, which in everyday life are often implicated in further important questions about how people will behave, how they ought to behave, the extent to which we credit or blame them for the outcome, and whether we excuse or punish them for transgressions.

We made three main findings, each replicated across different cover stories. First, we found that knowledge

attributions are insensitive to the luck of a salient threat that fails to prevent someone from perceptually detecting the truth. For example, suppose that a shopper purchases a diamond, puts it in her pocket, and walks out of the store. People overwhelmingly judge that the shopper knows that she has a diamond in her pocket as she leaves the store, even if a skilled pickpocket almost stole the diamond from her pocket before she left the store.

Second, we found that knowledge attributions are sensitive to the luck involved with an unnoticed disruption and change in the explanation for why a belief is true. For example, consider our shopper as she walks out of the store. If the diamond is in her pocket because she put it there, then people overwhelmingly judge that she knows that a diamond is in her pocket. However, if a pickpocket steals the diamond she originally put in her pocket, but her belief is nevertheless true because someone also secretly slipped a diamond into her pocket, then people are much less inclined to attribute knowledge to her.

Third, we found that when the explanation changes for why a belief is true, knowledge attributions are sensitive to the way in which the truth is replaced or restored. People are more inclined to attribute knowledge when the “backup” or replacement truth-maker is similar to the original truth-maker than when it is dissimilar. For example, consider our shopper, who believes that she has a diamond in her coat pocket, but whose diamond is stolen before she leaves the store. People are more inclined to attribute knowledge if her belief is true because the thief felt guilty and slipped the diamond back into her pocket than if it is true because her grandmother long ago sewed a diamond into the coat pocket. Although we have treated the similarity/dissimilarity distinction as a dichotomy in these initial attempts to test its effect, we acknowledge that it is probably better thought of as a continuum. Further work will be needed to understand how subtler changes in similarity affect knowledge attributions. Moreover, things can be similar in one respect but different in another. For instance, a male and a female squirrel are similar insofar as they are squirrels, but they differ in their genders. Further work is also needed to more precisely identify which sorts of similarity and dissimilarity affect knowledge attributions.

Though comparing results from different experiments is fraught, it is still worth noting the impressive consistency of knowledge attributions in structurally analogous conditions across Experiments 1–4. All knowledge controls consistently resulted in approximately 80 %–90 % knowledge attribution. Conditions involving failed threats were consistently treated similarly to clear cases of knowledge. The cases serving as ignorance controls consistently resulted in approximately 10 %–15 % knowledge attribution. And the cases of luck involving similar replacement consistently resulted in approximately 50 %–60 % knowledge attribution. The similarities among these findings suggest that the ordinary concept of

knowledge is deeply sensitive to the structural features of cases that we have identified.

A recent controversy in psychology and philosophy has concerned whether laypeople attribute lucky knowledge to agents in an intriguing range of cases known as “Gettier cases” (Gettier, 1963; see Turri, 2012, for a review). Studying Gettier cases is theoretically useful because they provide an excellent “opportunity to test which factors affect knowledge attributions” (Starmans & Friedman, 2013, p. 663). Some researchers have found that laypeople do tend to attribute knowledge in some Gettier cases (Starmans & Friedman, 2012), some researchers claim that laypeople do not (Nagel et al., 2013b), and others have observed mixed results, depending on the method of questioning (Turri, 2013). Our findings suggest an explanation for the seemingly inconsistent prior findings and theorizing on Gettier cases: Knowledge attributions are sensitive to different forms and combinations of luck, and prior research on Gettier cases has not controlled for all of the sensitivities identified here. Indeed, according to some criteria (e.g., Lewis, 1996; Pritchard, 2005; Zagzebski, 1996), researchers would count the stories used for five of the separate conditions in Experiment 4 as Gettier cases: Failed Threat, Detect Similar, Detect Dissimilar, No Detect Similar, and No Detect Dissimilar. But if intuitions about Gettier cases vary as widely as our results indicate—from patterns that closely resemble responses to *paradigmatic knowledge* (Failed Threat) to patterns that closely resemble responses to *paradigmatic ignorance* (no detect dissimilar)—then “Gettier case” is a theoretically useless category. The fact that something is a Gettier case would be 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*).

Starmans and Friedman (2012) identified an important distinction between cases in which the agent bases her belief on “authentic” evidence and cases in which she believes it on the basis of “apparent” evidence. They found that people are more likely to ascribe knowledge when the evidence was authentic than when it was merely apparent. Their cases of authentic evidence were structurally similar to several stories used in our experiments: Replacement (Exp. 2), Similar (Exp. 3), and Detect Similar (Exp. 4). Their cases of apparent evidence roughly corresponded to the stories used in No Detection (Exp. 2) and No Detect Similar (Exp. 4). Several of our results corroborate Starmans and Friedman’s (2012) findings, further supporting their view that the ordinary concept of knowledge is sensitive to the authentic/apparent distinction. The three most relevant results are that knowledge attribution was (i) higher in the Replacement than in the No Detection condition of Experiment 2, (ii) higher than chance rates in the Replacement condition of Experiment 2, and (iii)

higher in the Detect Similar than in the No Detect Similar condition of Experiment 4. It is also worth noting that the results from Experiment 2 generalize Starmans and Friedman's (2012) findings to cases that focus on natural kinds rather than on artifacts and to cases that do not involve stealing or tampering with people's property. In one experiment, we also observed knowledge attribution rise significantly above chance rates in an authentic-evidence case (Replacement, Exp. 2), which replicates a result reported by Starmans and Friedman (2012). Nevertheless, we also observed that this high rate of knowledge attribution does not always replicate (Similar, Exp. 3), which suggests that still unidentified factors affect knowledge attribution in such cases.

Beyond the question of how much above chance the true average for the population is in such cases, the present results lend further strong support to Starmans and Friedman's (2012) conclusion that philosophers have been too quick to claim that all Gettier cases are obvious cases of ignorance. Nagel and colleagues (Nagel et al., 2013a, b) have recently empirically challenged Starmans and Friedman's (2012) conclusions, arguing that laypeople view apparent-evidence cases as cases of ignorance. Starmans and Friedman (2013) correctly noted that the challenge was partly based on a critical statistical miscalculation. Moreover, the present findings on authentic-evidence cases contradict the empirical claims made by Nagel and colleagues.

In the philosophical study of knowledge, patterns in ordinary knowledge attributions are typically taken to provide an important constraint on theories of knowledge. Analytic epistemologists have largely followed J. L. Austin, a principal founder of "ordinary language" language philosophy, when he wrote,

If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing. . . . Certainly ordinary language is *not* the last word: in principle it can everywhere be supplemented and improved upon and superseded. Only remember, it is the *first* word. (Austin, 1956, p. 11; cf. Reid, 1785/2002, pp. 26–27).

More recently, leading epistemologists have also emphasized a strong preference for respecting patterns in ordinary judgments about knowledge. According to Jonathan Vogel, a theory of knowledge should "accommodate the body of our intuitions in an unforced, convincing way" (Vogel, 1990, p. 298). According to John Hawthorne, a criterion of a good theory of knowledge is that it respects "uncontroversial" and "widespread" tendencies to ascribe or deny knowledge in certain cases (Hawthorne, 2004, chap. 1 et passim; see also Stroud, 1984, and DeRose, 2009). Given the premium that contemporary epistemologists place on fidelity to ordinary usage, the present study, and other psychological studies of

knowledge attribution, can helpfully inform substantive philosophical theorizing about knowledge, especially theorizing about "epistemic luck," or the types of luck that are inconsistent with knowledge.

An important question for further research is why knowledge attributions are selectively sensitive to some but not to other forms of luck. Answering this question might also help shed light on the social function of knowledge attributions, as well as on neighboring domains such as *moral* evaluation. Ethicists and moral psychologists have asked similar questions about the role of luck in evaluating an agent's moral status (Williams & Nagel, 1976; Young, Nichols, & Saxe, 2010). But, whatever the ultimate explanation for this particular pattern of selective sensitivity, one thing is certain. We are fallible social beings deeply interested in keeping track of what people deserve and are responsible for. Whether it is a question of knowledge or morality or athletics, we seek criteria to distinguish genuine achievements from lucky outcomes. Our results have helped to expose some of those criteria.

Acknowledgements This research was supported by the Social Sciences and Humanities Research Council of Canada, The Character Project at Wake Forest University and the John Templeton Foundation (neither of which necessarily endorses any opinion expressed here), and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

References

- Aristotle. (1941). *Posterior analytics*. In R. McKeon (Ed.), G. R. G. Mure (Trans.), *The basic works of Aristotle*. New York, NY: Random House
- Austin, J. L. (1956). A plea for excuses. *Proceedings of the Aristotelian Society*, 57, 1–30.
- Beebe, J. R., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind and Language*, 25, 1–25.
- Beebe, J. R., & Shea, J. (2013). Gettierized Knobe effects. *Episteme*, 10, 219–240. doi:10.1017/epi.2013.23
- Blouw, P., Buckwalter, W., & Turri, J. (in press). *Gettier cases: a taxonomy*. In Borges, R., de Almeida, C., & Klein, P (Eds.), *Explaining knowledge: new essays on the Gettier problem*. Oxford: Oxford University Press.
- Booth, J. R., Hall, W. S., Robison, G. C., & Kim, S. Y. (1997). Acquisition of the mental state verb know by 2- to 5-year-old children. *Journal of Psycholinguistic Research*, 26, 581–603. doi:10.1023/A:1025093906884
- Buckwalter, W. (2013). Gettier made ESEE. *Philosophical Psychology*, 27, 368–383. doi:10.1080/09515089.2012.730965
- Buckwalter, W. (in press). *Factive verbs and protagonist projection*. Episteme.
- Buckwalter, W., & Turri, J. (in press). Telling, showing and knowing: A unified account of pedagogical norms. *Analysis*.
- Chisholm, R. (1989). *Theory of knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Colaço, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11, 199–212. doi:10.1017/epi.2014.7
- Cullen, S. (2010). Survey-driven Romanticism. *Review of Philosophy and Psychology*, 1, 275–296. doi:10.1007/s13164-009-0016-1

- DeRose, K. (2009). *The case for contextualism*. Oxford, UK: Oxford University Press.
- Duff, R. A. (1990). *Intention, agency and criminal liability*. Malden, MA: Blackwell.
- Engel, M., Jr. (1992). Is epistemic luck compatible with knowledge? *Southern Journal of Philosophy*, 30, 59–75. doi:10.1111/j.2041-6962.1992.tb01715.x
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford, UK: Oxford University Press.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39, 175–191. doi:10.3758/BF03193146
- Feltz, A., & Cokely, E. T. (2012). The philosophical personality argument. *Philosophical Studies*, 161, 227–246. doi:10.1007/s11098-011-9731-4
- Friedman, O., & Turri, J. (in press). Is probabilistic evidence a source of knowledge? *Cognitive Science*.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge, UK: Cambridge University Press.
- Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.
- Hart, H. (1959). Prolegomenon to the principles of punishment. *Proceedings of the Aristotelian Society*, 60, 1–26.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford, NY: Oxford University Press.
- Hazlett, A. (2010). The myth of factive verbs. *Philosophy and Phenomenological Research*, 80, 497–522.
- Hetherington, S. (1998). Actually knowing. *Philosophical Quarterly*, 48, 453–469. doi:10.1111/1467-9213.00114
- Hetherington, S. (1999). Knowing failably. *Journal of Philosophy*, 96, 565.
- Hetherington, S. (2013). Knowledge can be lucky. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 164–176). Malden, MA: Wiley-Blackwell.
- Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16, 309–324. doi:10.1080/09515080307771
- Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25, 168–175.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567. doi:10.1080/00048409612347521
- Lycan, W. G. (2006). On the Gettier Problem problem. In S. Hetherington (Ed.), *Epistemology futures* (pp. 148–168). Oxford, UK: Oxford University Press.
- Malle, B. (2006). Intentionality, morality, and their relationship in human judgment. *Journal of Cognition and Culture*, 6, 87–112. doi:10.1163/156853706776931358
- Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33, 101–121.
- Malle, B. F., & Nelson, S. E. (2003). Judging mens rea: The tension between folk concepts and legal concepts of intentionality. *Behavioral Sciences & the Law*, 21, 563–580. doi:10.1002/bsl.554
- Mele, A. R., & Cushman, F. (2007). Intentional action, folk judgments, and stories: Sorting things out. *Midwest Studies in Philosophy*, 31, 184–201. doi:10.1111/j.1475-4975.2007.00147.x
- Nagel, J., Mar, R., & San Juan, V. (2013a). Authentic Gettier cases: A reply to Starmans and Friedman. *Cognition*, 129, 666–669. doi:10.1016/j.cognition.2013.08.016
- Nagel, J., San Juan, V., & Mar, R. A. (2013b). Lay denial of knowledge for justified true beliefs. *Cognition*, 129, 652–661. doi:10.1016/j.cognition.2013.02.008
- Pizarro, D. A., Uhlmann, E., & Bloom, P. (2003). Causal deviance and the attribution of moral responsibility. *Journal of Experimental Social Psychology*, 39, 653–660.
- Plato. (1992). *The republic* (G. M. A. Grube & C. D. C. Reeve, Trans.). Indianapolis, IN: Hackett.
- Pritchard, D. (2005). *Epistemic luck*. New York, NY: Oxford University Press.
- Reid, T. (2002). In D. Brookes (Ed.), *Essays on the intellectual powers of man*. University Park, PA: Pennsylvania State University Press. Original work published 1785.
- Russell, B. (1948). *Human knowledge: Its scope and limits*. New York, NY: Routledge.
- Sartwell, C. (1991). Knowledge is merely true belief. *American Philosophical Quarterly*, 28, 157–165.
- Sartwell, C. (1992). Why knowledge is merely true belief. *Journal of Philosophy*, 89, 167–180.
- Searle, J. (1983). *Intentionality*. Cambridge, UK: Cambridge University Press.
- Shatz, M., Wellman, H. M., & Silber, S. (1983). The acquisition of mental verbs: A systematic investigation of the first reference to mental state. *Cognition*, 14, 301–321.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*. Princeton, NJ: Princeton University Press.
- Shope, R. K. (2002). Conditions and analyses of knowing. In P. K. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 25–70). Oxford, UK: Oxford University Press.
- Sodian, B. (1988). Children's attributions of knowledge to the listener in a referential communication task. *Child Development*, 59, 378–385.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge (Vol. I)*. Oxford, UK: Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford, UK: Oxford University Press.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124, 272–283.
- Starmans, C., & Friedman, O. (2013). Taking “know” for an answer: A reply to Nagel, San Juan, and Mar. *Cognition*, 129, 662–665. doi:10.1016/j.cognition.2013.05.009
- Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford, UK: Oxford University Press, Clarendon Press.
- Turri, J. (2011a). The express knowledge account of assertion. *Australasian Journal of Philosophy*, 89, 37–45.
- Turri, J. (2011b). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, 11(8), 1–11.
- Turri, J. (2011c). Mythology of the factive. *Logos and Episteme*, 2, 143–152.
- Turri, J. (2012). In Gettier's wake. In S. Hetherington (Ed.), *Epistemology: The key thinkers* (pp. 214–229). London, UK: Continuum.
- Turri, J. (2013). A conspicuous art: Putting Gettier to the test. *Philosophers' Imprint*, 13(10), 1–16.
- Turri, J. (2014a). *Epistemology: A guide*. Malden, MA: Wiley-Blackwell.
- Turri, J. (2014b). Knowledge and supererogatory assertion. *Philosophical Studies*, 167, 557–567. doi:10.1007/s11098-013-0112-z
- Turri, J., & Buckwalter, W. (2014). Descartes's schism, Locke's reunion: completing the pragmatic turn in epistemology. Manuscript under review.
- Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, 65, 157–170.
- Unger, P. (1975). *Ignorance: A case for skepticism*. Oxford, UK: Oxford University Press.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting: Contemporary perspectives on skepticism* (pp. 13–29). Dordrecht, The Netherlands: Kluwer. (Reprinted from *Epistemology: An Anthology*, 2nd ed, pp. 290–301, by E. Sosa, J. Kim, J. Fantl, & M. McGrath, Eds., 2008, Wiley-Blackwell). doi:10.1007/978-94-009-1942-6_2

- Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, *29*, 429–460.
- Williams, B., & Nagel, T. (1976). Moral luck. *Proceedings of the Aristotelian Society*, *50*, 115–151.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford, UK: Oxford University Press.
- Wittgenstein, L. (1975). *On certainty*. In: Paul & G. Anscombe, Trans.; G. Anscombe & G. H. V. Wright (Eds.). Malden, MA: Blackwell.
- Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, *115*, 491–503. doi:[10.1016/j.cognition.2010.02.003](https://doi.org/10.1016/j.cognition.2010.02.003)
- Young, L., Nichols, S., & Saxe, R. (2010). Investigating the neural and cognitive basis of moral luck. *Review of Philosophy and Psychology*, *1*, 333–349.
- Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge, UK: Cambridge University Press.