



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

The folk conception of knowledge

Christina Starmans^{a,*}, Ori Friedman^{b,*}^a Department of Psychology, 2 Hillhouse Ave., New Haven, CT 06511, USA^b Department of Psychology, University of Waterloo, 200 University Ave. West, Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:

Received 23 December 2011

Revised 22 May 2012

Accepted 23 May 2012

Available online xxxx

Keywords:

Gettier cases

Knowledge

Theory of mind

Folk psychology

Epistemic intuitions

ABSTRACT

How do people decide which claims should be considered mere beliefs and which count as knowledge? Although little is known about how people attribute knowledge to others, philosophical debate about the nature of knowledge may provide a starting point. Traditionally, a belief that is both true and justified was thought to constitute knowledge. However, philosophers now agree that this account is inadequate, due largely to a class of counterexamples (termed “Gettier cases”) in which a person’s justified belief is true, but only due to luck. We report four experiments examining the effect of truth, justification, and “Gettiering” on people’s knowledge attributions. These experiments show that: (1) people attribute knowledge to others only when their beliefs are both true and justified; (2) in contrast to contemporary philosophers, people also attribute knowledge to others in Gettier situations; and (3) knowledge is not attributed in one class of Gettier cases, but only because the agent’s belief is based on “apparent” evidence. These findings suggest that the lay concept of knowledge is roughly consistent with the traditional account of knowledge as justified true belief, and also point to a major difference between the epistemic intuitions of laypeople and those of philosophers.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

People espouse all kinds of incredible beliefs: The world will end on May 21, 2011; aliens landed in Roswell, New Mexico; invisible entities called germs cause us to get sick; the universe is composed of eleven dimensions of vibrating strings. To assess these claims, we need to decide which we should regard as mere beliefs or opinions, and which we should treat as *knowledge*. But what is knowledge, and what is required for a fact to be known?

These questions have traditionally been viewed as questions for philosophers. However, ordinary people make frequent decisions about what others know. How people make these knowledge attributions is a psychological question, and one about which little is known. This is especially

surprising given the vast number of studies investigating people’s “theory of mind” reasoning; while researchers have developed detailed accounts of how people attribute beliefs (e.g., Apperly & Butterfill, 2009; Leslie, Friedman, & German, 2004; Nichols & Stich, 2003), little is known about how people attribute knowledge.

Recently, many researchers have examined how people evaluate *who* is a good source of knowledge. This research shows that even children as young as 3–4 years do not blindly accept the testimony of others as fact, but selectively trust speakers that have been accurate in past, as well as speakers that are older, more familiar, more certain, and more expert (e.g., Birch, Vauthier, & Bloom, 2008; Corriveau & Harris, 2009; Jaswal & Neely, 2006; Koenig, Clement, & Harris, 2004; Lutz & Keil, 2002; Sabbagh & Baldwin, 2001; Taylor, Cartwright, & Bowden, 1991). Earlier research directly examining knowledge attributions investigated the ages at which children infer that access to different types of information leads to knowledge or ignorance (e.g., Pratt & Bryant, 1990; Sodian, 1988; Sodian & Wimmer, 1987;

* Corresponding authors. Tel.: +1 203 432 0877 (C. Starmans), tel.: +1 519 888 4567x33054 (O. Friedman).

E-mail addresses: christina.starmans@yale.edu (C. Starmans), friedman@uwaterloo.ca (O. Friedman).

Wimmer, Hogrefe, & Perner, 1988; Woolley & Wellman, 1993); and also children's difficulty ignoring their own privileged knowledge when reasoning about whether others are knowledgeable (Birch & Bloom, 2007). However, this research has not provided an account of *how* people attribute knowledge.

Although little is known about how people normally attribute knowledge, we do know quite a bit about how one group of people think about knowledge: Philosophers have inquired into the nature of knowledge for millennia. Plato (*Meno*, *Theatetus*) described knowledge as the possession of a justified true belief. Suppose Sam is at the race-track, and believes that his favorite horse, Lucky, has a silver dollar under her saddle because he saw Lucky's jockey put one there. According to this view of knowledge, Sam knows there is a silver dollar under Lucky's saddle because he believes there is a silver dollar under Lucky's saddle, and because this belief is both true (there is a silver dollar under her saddle) and justified (Sam has a good reason to believe there is a silver dollar under Lucky's saddle because he saw her jockey put one there). This view of knowledge has intuitive appeal. If Sam claims that he knows that a silver dollar is under Lucky's saddle, he implies (1) that the silver dollar is there, and (2) that his assertion that it is there can be trusted. Hence his claim of knowledge would be contradicted if his belief were false (e.g., the coin under Lucky's saddle is actually a quarter), and it would be contradicted if his belief were unjustified (e.g., Sam never saw the coin, and only formed his belief because of a wild guess). If Sam's belief were false or unjustified, it would be just another thing that Sam believes but does not really know.

Although the JTB view of knowledge is appealing, philosophers now view it as inadequate. Edmund Gettier (1963) pointed out that there are situations in which a person has a belief that is both justified and true, but would probably not be considered to have knowledge. Returning to Sam, suppose that one day at the track Sam sneaks into Lucky's stall minutes before the race and tucks a silver dollar under her saddle. However, a rival jockey sees him, and steals the silver dollar after Sam leaves. Happily though, Lucky's jockey finds a different silver dollar and slips it under Lucky's saddle before the race. As he watches the race, Sam believes there is a silver dollar under Lucky's saddle, and this belief is both justified and true. Nonetheless, it is commonly claimed that the overwhelming consensus among philosophers is that Sam does not know that there is a silver dollar under Lucky's saddle (e.g., Lycan, 2006; Turri, 2012; Williamson, 2005).¹

Philosophers have never agreed on a precise account of why Sam lacks knowledge when he is "Gettiered". However, many proposed explanations center on two properties shared by Gettier cases. One property is the element of luck (Pritchard, 2005; Unger, 1968): Sam does not know there is a silver dollar under Lucky's saddle because his belief is only true due to luck, or more precisely, due to "double luck" (Turri, 2011; Zagzebski, 1994). Though Sam

starts with a justified true belief, bad luck interferes (the silver dollar is stolen), and then a stroke of good luck "cancels out the bad" (a different silver dollar ends up under the saddle).

The second property shared by Gettier cases is that the fact that justifies the belief (Sam put a silver dollar under Lucky's saddle) is causally disconnected (Goldman, 1967) from the fact that makes the belief true (the jockey put a silver dollar under Lucky's saddle). On this view, Sam does not know there is a silver dollar under Lucky's saddle because his belief, though true and justified, is true for the wrong reason.

Regardless of which account of Gettier intuitions is correct, examining laypeople's intuitions about these cases will provide important insights into how people ordinarily attribute knowledge. If laypeople share philosophers' intuitions that a Gettiered individual lacks knowledge, this would suggest that everyday attributions of knowledge depend on causal reasoning or on detecting luck. However, if people instead view Gettiered individuals as possessing knowledge, this would suggest that their knowledge attributions do not depend on such reasoning, and could indicate that people instead reason according to the classic view that knowledge is justified true belief.

Previous investigations by experimental philosophers have examined people's intuitions about Gettier cases with varying results. For example, Weinberg, Nichols, and Stich (2001) tested college students from a variety of cultural backgrounds using a story about a man named Bob who believed his coworker Jill owned an American car because he had seen her driving a Buick. However, Jill's Buick had been recently stolen, and she had replaced it with a Pontiac—another American car. Participants were asked whether Bob *really knows* Jill owns an American car, or whether he *only believes* it. Only 26% of Westerners attributed knowledge in this Gettier case, but a substantial number of Asians and Indians (53% and 61%, respectively) attributed knowledge. Cullen (2010) repeated this study but asked simply whether Bob *knows* or *does not know* that Jill owns an American car, and found that 42% of his North American participants now attributed knowledge to Bob. Buckwalter (2012) describes a study in which participants read a Gettier case about a CEO signing some documents, and found that when asked to attribute knowledge on a Likert scale, participants scored above the midpoint, suggesting that they considered this Gettier case to constitute knowledge.

However, a common problem unites all of these experiments. The Gettier cases presented to participants were not compared with well-structured control cases, and so none of these experiments demonstrate that participants' responses were related to the Gettiering in the scenario. For example, in Weinberg et al.'s scenario about the American car, it is possible that their Western participants would have been skeptical across the board, denying that Bob knew that Jill owned an American car even if Jill's car had not been stolen (for example, in a closely matched non-Gettier control in which a *different* car was stolen). Similarly, perhaps the participants who attributed knowledge in these scenarios would also attribute knowledge in cases where beliefs were unjustified and/or untrue—if, for

¹ While this claim has not been empirically tested, for ease of exposition we will follow the vast majority of the philosophical literature in assuming this intuition is in fact shared by most philosophers.

example, people are inclined to attribute knowledge to someone who is simply very certain. Thus, to carefully examine whether people attribute knowledge in Gettier cases, it is important to demonstrate that they attribute knowledge in closely matched control scenarios where the character has a justified true belief and is not Gettiered, and that they do not attribute knowledge in closely matched scenarios in which the belief is either not true or not justified.

Such comparisons form the basis of the current paper. We provide evidence that people attribute knowledge in cases where a person possesses a justified true belief, and do so even when the person is Gettiered. Broadly consistent with the idea that people view knowledge as justified true belief, these experiments also show that people do not attribute knowledge to individuals with false beliefs, nor to those possessing unjustified true beliefs. Finally, we also identify a class of Gettier cases where people do not attribute knowledge. However, we demonstrate that this is not because the agent is Gettiered, but rather because the agent's belief is based on apparent evidence.

2. Experiment 1A

2.1. Method

2.1.1. Participants

One hundred and forty-four participants (72 female, aged 18–81 yrs, mean = 31 yrs, standard deviation = 12 yrs) were recruited using Amazon Mechanical Turk (<http://www.mturk.com/>). Participants were located throughout the United States, and 98% listed English as their native language. Education levels ranged from some high school to a PhD or professional degree; 70% selected either "Some College" or Bachelor's Degree". Only 19% had taken more than one course in philosophy. Participants were paid \$0.20 for approximately 1–2 min of time. To screen out repeat participants, we asked that people not participate if they had previously taken a similar survey, and we eliminated data from participants with identical Worker IDs and/or IP addresses.² Data were also excluded from a further 38 participants who failed comprehension questions designed to ensure participants were reading the scenarios carefully; however, including these participants yielded the same pattern of results as described below.

2.1.2. Materials and procedure

Participants were randomly assigned to either a Gettier, Control, or False Belief condition, and each read the following scenario (text from the False Belief condition in brackets):

Peter is in his locked apartment reading, and is about to have a shower. He puts his book down on the coffee table, and takes off his black plastic watch and leaves it on the coffee table. Then he goes into the bathroom.

² Participants for subsequent experiments were also recruited using Amazon Mechanical Turk. Repeat participants were screened out in the same fashion (we excluded anyone who had participated in any of the previous studies); demographics were similar across all experiments.

As Peter's shower begins, a burglar silently breaks into the apartment. The burglar takes Peter's black plastic watch, replaces it with an identical black plastic watch [a dollar bill], and then leaves. Peter is still in the shower, and did not hear anything.

Following the story, participants answered a series of comprehension questions designed to ensure they were attending to the relevant elements of the story, followed by a knowledge question and a confidence question. In the Gettier and False Belief conditions these questions referred to an object that had been switched (a watch); in the Control condition the questions referred to an object that had not been switched (a book). The questions were as follows; response options are in brackets.

At the end of the story:

1. Is there a watch [book] on the table? (Yes/No)
2. How did the watch [book] get on the table? (Peter put it there/The burglar put it there)
3. Would Peter say there is a watch [book] on the table?³ (Yes/No)
4. Why would Peter say there is a watch [book] on the table? (Because Peter put a watch [book] on the table/ Because a burglar put a watch [book] on the table)
5. Peter _____ that there is a watch on the table. (Really knows/Only thinks)
6. How confident are you about your answer to Question 5 (above)? (1 – Not at all confident to 10 – Completely confident)

In this and all subsequent experiments, questions were always presented in the same order, while response options were presented in random order for all questions except the confidence Likert scale. Each question appeared on a new screen, except Questions 5 and 6 were shown together. Participants could not go back to alter a previous answer, and the story remained at the top of the screen throughout. Following testing, participants filled out a brief demographic questionnaire; this questionnaire also followed testing in the subsequent experiments.

In the Gettier condition, Peter has a justified true belief that there is a watch on the table (because he left a watch on the table, and there is a watch on the table). However, he has also been Gettiered (because the watch he left on the table is not the watch that is now on the table). Hence, if participants share the intuitions of philosophers, they should judge that Peter is not knowledgeable in this condition. If participants instead view knowledge as justified true belief, they should attribute knowledge in this condition.

³ Questions 3 and 4 were intended to make participants reflect on what Peter believes about the watch or book. However, pilot testing in previous experiments indicated that many laypeople interpret "thinks" or "believes" in opposition to "knows". Thus participants often indicated that an agent does not believe something, but does know it—and in fact that he does not believe it because he knows it. Therefore, in questions 3 and 4 we asked what Peter would say as a proxy for what he believes. Although in some cases what an agent would say differs from what he believes, we did not find a greater proportion of errors on these questions than on any other questions.




Fig. 1. Experiment 1A. Left panel: Mean knowledge scores derived by combining dichotomous choice with confidence ratings. Right panel: Percentage of participants attributing knowledge in the Gettier, Control, and False Belief conditions.

In the Control condition, Peter has a justified true belief and is not Gettiered; he put a book on the table and the same book is still on the table. Thus we expected that most participants would attribute knowledge here. However, ceiling rates of knowledge were not necessarily expected—some participants might take the skeptical position that one can never *really know* anything, and thus decline to attribute knowledge in any condition.

Lastly, we expected that participants would decline to attribute knowledge in the False Belief condition because Peter's belief, though justified, is not true—although he believes there is a watch on the table, it was replaced by something else. This condition is important because it allows us to confirm that participants are not interpreting “knows” as equivalent to a confidently-held belief, and that they instead appreciate that knowledge depends not just on one's mental states but also on the state of the world.

2.2. Results and discussion

If participants share the intuitions of philosophers, they should attribute knowledge more in the Control condition than in the Gettier and False Belief conditions. However, if people instead view knowledge as justified true belief, they should attribute knowledge similarly in the Gettier and Control conditions, but less in the False Belief condition. To test between these views, we combined the two measures used to assess knowledge judgments (i.e., the dichotomous choice between ‘really knows’ and ‘only believes’, and the 1-to-10 confidence scale), multiplying the confidence rating by +1 in cases where participants attributed knowledge, and by -1 in cases where participants did not attribute knowledge. Thus each participant's score fell on a 20-point scale ranging from -10 (minimum knowledge attribution) to +10 (maximum knowledge attribution).

A one-way analysis of variance (ANOVA) revealed an overall effect of condition ($F(2, 141) = 63.65, p < .001$).⁴ As shown in Fig. 1, knowledge attribution exceeded chance in both the Gettier ($t(46) = 4.00, p < .001$, all tests two-tailed)

and Control conditions ($t(50) = 8.75, p < .001$), though participants attributed knowledge to Peter more often in the Control condition ($M = 7.27, SD = 5.94$) than in the Gettier condition ($M = 4.53, SD = 7.78, t(96) = -1.97, p = .05$). In the False Belief condition ($M = -7.13, SD = 5.97$), knowledge was attributed less than in the Gettier condition ($t(91) = 8.10, p < .001$), and at rates less than would be expected by chance ($t(45) = -8.11, p < .001$).

A similar pattern emerged when only examining responses to the dichotomous knowledge question: participants attributed knowledge at rates greater than chance in both the Control condition (88%) and the Gettier condition (72%, binomial, both $ps < .003$), but here knowledge attribution did not differ between the two conditions (Fisher's exact test, $p = .07$, see Fig. 1). Knowledge was attributed less in the False Belief (11%) than in the Gettier condition (Fisher's, $p < .001$), and less than would be expected by chance (binomial, $p < .001$). Confidence levels did not differ across conditions ($M = 9.1/10, F(6, 120) = 1.18, p = .32$).

Participants attributed knowledge to a person who had an uncontroversial justified true belief and to a person who was Gettiered, and did not attribute knowledge to a person with a false belief. This reflects the pattern expected if people generally think of knowledge as justified true belief, and conflicts with the outcomes expected if they share the intuitions of philosophers.

The Gettier scenario that participants read in this experiment features the typical causal disconnect between the fact justifying Peter's belief (Peter left his watch on the table) and the fact making the belief true (the thief left a watch on the table). However, one might point out that there is *some* causal connection between the fact justifying the belief and the fact making the belief true, given that the thief might have put the second watch down *because* he took the original watch. Similarly, the “bad luck” and “good luck” are connected (because the thief who undermined the truth of the belief was also responsible for ensuring it remained true), and so one might question whether this constitutes “double luck”, or merely a single instance of luck.

However, the degree of double luck and causal disconnect in this story is typical of many scenarios counted as “Gettier cases”, and for which philosophers have declined to attribute knowledge (e.g., “Contract”, Buckwalter, 2012;

⁴ Initial omnibus ANOVAs showed no effect of sex, education level, or order of presentation on knowledge attribution scores for this or any subsequent experiment.

“Hologram of man”, Fantl & McGrath, 2009; “Burgled”, Sturgeon, 1993; “Hobbled”, Turri, 2011; “Special dog”, Turri, 2012; “American car”, Weinberg et al., 2001; “Candles on seesaw”, Williams, 1978). For example, the philosopher Scott Sturgeon (1993, p. 157) offers the following as “a typical Gettier scenario”:

Suppose I burgle your house, find two bottles of Newcastle Brown in the kitchen, drink and replace them. You remember purchasing the ale and come to believe there will be two bottles waiting for you at home. Your belief is justified and true, but you do not know what’s going on.

This scenario is extremely similar to the Gettier case our participants considered, and also features a “replacing thief”. Nonetheless (and setting aside debate about what counts as a Gettier case), participants might decline to attribute knowledge if they considered a Gettier scenario in which the “good luck” and “bad luck” are not connected, and in which the fact justifying the belief is more thoroughly disconnected from the fact making the belief true. The next experiment explored this possibility.

3. Experiment 1B

3.1. Method

3.1.1. Participants

One hundred thirty-three participants (67 female, aged 18–65 yrs, mean = 31 yrs, standard deviation = 10.5 yrs) were tested. Data were excluded from an additional 35 participants who failed comprehension questions; again, including these participants yielded no significant differences in the pattern of results.

3.1.2. Materials and procedure

Participants were randomly assigned to one of four conditions. In each condition, participants read a modified version of the story from Experiment 1A, and then answered the same 6 questions as in Experiment 1A. Participants in the 1-Thief Gettier condition read a story in which a burglar steals a pen and replaces it with another pen. This story was very similar to the Gettier condition in Experiment 1A, and was included as a replication. Participants in the remaining three conditions (2-Thief Gettier, Control, and False Belief) read scenarios involving two thieves. In the Gettier version, one thief steals a pen, but the other thief accidentally leaves an identical pen in its place. In this story, the bad luck interfering with the truth of the agent’s belief is completely disconnected from the good luck that makes the belief true; moreover there is no causal connection between the fact justifying the belief and the fact making the belief true. Here is the 2-Thief scenario (text from the False Belief condition is in brackets):

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 absentmindedly

leaves his own identical blue Bic pen [his own bandana] on the coffee table. Then the burglars leave. Katie is still in the shower, and did not hear anything.

3.2. Results and discussion

Replicating Experiment 1A, knowledge attribution exceeded chance in the 1-Thief Gettier condition ($M = 3.77$, $SD = 8.42$, $t(38) = 2.79$, $p = .008$). Participants also attributed knowledge at rates exceeding chance in the 2-Thief Gettier condition ($M = 3.92$, $SD = 8.21$, $t(25) = 2.44$, $p = .02$), and importantly, rates of knowledge attribution did not differ between these two Gettier conditions ($t(63) = -.07$, $p = .94$, see Fig. 2).

Comparing the three 2-Thief conditions reveals a similar pattern as in Experiment 1A. A one-way ANOVA again revealed an overall effect of condition ($F(2,91) = 23.75$, $p < .001$), with participants attributing knowledge to Katie equally in the Control condition ($M = 5.41$, $SD = 7.68$) and in the Gettier condition ($M = 3.92$, $SD = 8.21$, $t(63) = -.74$, $p = .46$), but less in the False Belief condition ($M = -6.69$, $SD = 6.62$) than in the Gettier condition ($t(66) = -6.81$, $p < .001$). Knowledge was attributed at rates exceeding chance in both the Gettier and Control conditions ($t(38) = 4.40$, $p < .001$), and at rates less than would be expected by chance in the False Belief condition ($t(45) = -8.13$, $p < .01$).

Similarly, when only examining responses to the dichotomous knowledge question participants were equally likely to attribute knowledge in the 1-Thief (72%) and 2-Thief Gettier conditions (69%, Fisher’s, $p = 1.0$, see Fig. 2). Participants also attributed knowledge equally in the 2-Thief Gettier condition and the 2-Thief Control condition (79%, Fisher’s, $p = .39$), and did so at rates greater than chance in all three Gettier and Control conditions (binomial, all $ps < .04$). Knowledge was attributed less in the 2-Thief False Belief (14%) than in the 2-Thief Gettier condition (Fisher’s, $p < .001$), and less than would be expected by chance (binomial, $p < .001$). Confidence levels did not differ across conditions ($M = 9.1/10$, $F(3, 129) = .46$, $p = .71$).

Experiment 1B thus replicated the main finding from Experiment 1A, as participants again attributed knowledge to a person who had an uncontroversial justified true belief and to a person who was Gettiered, and did not attribute knowledge to a person with a false belief. Moreover, this experiment also finds that participants do not distinguish between a Gettier scenario in which the “double luck” elements are connected and a scenario in which they are not connected. Likewise, highlighting and extending the causal disconnect between the fact justifying the belief and the fact making the belief true did not affect participants’ judgment that the Gettiered agent possesses knowledge.

In both Experiments 1A and 1B participants were asked a series of comprehension questions relating to the protagonists’ belief, the actual state of the world, and the lack of causal connection between the reason for the belief and the reason that the believe is true. The rationale for these questions was twofold: (1) to ensure that our online participants were engaged and attending to the relevant parts of the scenarios; and (2) to closely approximate the logical steps that philosophers have deemed relevant for the




Fig. 2. Experiment 1B. Left panel: Mean knowledge scores derived by combining dichotomous choice with confidence ratings. Right panel: Percentage of participants attributing knowledge in the Gettier, Control, and False Belief conditions.

attribution of knowledge. Despite this extra guidance, participants attributed knowledge to characters who were Gettiered. However, perhaps asking this series of comprehension questions actually caused participants to attribute knowledge more than they would have otherwise, either because they were distracted or confused by the questions, or because the questions caused them to reflect on knowledge in ways that they would not naturally have done. To investigate this possibility, we gave a new group of participants the same scenario as in Experiment 1A, but this time asked only one comprehension question before asking whether the protagonist possessed knowledge.

4. Experiment 1C

4.1. Method

4.1.1. Participants

Forty-six participants (27 female, aged 18–57 yrs, mean = 32 yrs, standard deviation = 13 yrs) were tested. Data were excluded from an additional participant who failed the comprehension question.

4.1.2. Materials and procedure

Each participant read the Gettier scenario from Experiment 1A, and then answered only one comprehension question (“Is there a watch on the table?” [Yes/No]), the knowledge question, and the confidence question.

4.2. Results and discussion

Participants attributed knowledge to Peter at rates exceeding chance, both when examining the combined scores ($M = 5.74$, $t(45) = 5.93$, $p < .001$), and when only considering the dichotomous responses (83%, binomial, $p < .001$). These results do not differ from those found in any of the three Gettier conditions in the previous experiments (Fisher’s, all $ps > .24$). This suggests that regardless of whether participants are asked for an immediate response or forced to reason through the relevant factors in the scenario, laypeople nonetheless consider Gettier cases to be instances of knowledge.

The pattern of results observed so far is consistent with a lay concept of knowledge that requires beliefs to be

justified and true, but allows for justified beliefs to be true due to luck. However, these results are also consistent with a lay concept of knowledge as merely true belief. That is, perhaps laypeople do not require that beliefs be justified in order to count as knowledge, but only that they be true. To explore this, we next systematically manipulated whether the character’s belief was justified. In doing so we directly tested the prediction that justified true beliefs are viewed as knowledge, but that unjustified true beliefs are not.

5. Experiment 2

5.1. Method

5.1.1. Participants

Fifty-one new participants (31 female, aged 18–60 yrs, mean = 31 yrs, standard deviation = 10 yrs) were tested. Data were excluded from an additional three participants who failed control questions.

5.1.2. Materials and procedure

Participants were randomly assigned to one of two conditions, which differed in whether the protagonist left an object in a secure manner for a short period of time (High Justification) or in an insecure manner for a long period of time (Low Justification). Each participant read the following Gettier case; text in brackets was seen in the Low Justification condition.

Jane owns a small bookstore. One morning, while a few customers are browsing through the books, Jane puts a \$20 bill into the empty cash register. She locks it up [leaves it wide open] and goes into the back room to brew some coffee. While Jane is in the back room, her employee Bill arrives. He looks in the register and notices that the \$20 bill is a little worn. He takes it out of the cash register and replaces it with a new \$20 bill and leaves to run some errands. Jane has only been in the back room for a few minutes [has been in the back room for half an hour], and did not hear anything.

As in Experiment 1C, participants responded to one comprehension question, a knowledge question, and a confidence question.

5.2. Results and discussion

Once again a one-way ANOVA showed a main effect of condition ($F(1,49) = 11.75, p = .001$). Knowledge attribution exceeded chance in the High Justification condition ($M = 4.00, SD = 7.62, t(26) = 2.73, p = .01$), replicating Experiments 1A and 1B. In the Low Justification condition ($M = -3.29, SD = 7.54$), knowledge was attributed less than in the High Justification condition ($t(49) = 3.43, p = .001$, see Fig. 3), and at lower rates than would be expected by chance ($t(23) = -2.14, p = .04$). Participants were equally confident of their answers in both conditions ($M = 8.1/10, t(49) = .91, p = .37$).

A similar pattern emerged when only examining responses to the dichotomous knowledge question: participants attributed knowledge to Jane in the High Justification condition (70%, binomial, $p = .05$) but did not attribute knowledge to Jane when she was in a Low Justification situation (25%, binomial, $p = .02$, see Fig. 3). The difference between these conditions (Fisher's, $p = .002$) suggests that participants consider it necessary for a belief to be justified in order for it to count as knowledge.

The results of Experiments 1 and 2 suggest that people view justified true beliefs as knowledge even when agents are Gettiered. However, we next suggest that there is a subset of Gettier cases in which people do not attribute knowledge, though not because the character is Gettiered.

In many (though not all) of the Gettier cases frequently discussed by philosophers, an agent forms a belief based on “apparent” evidence—evidence that only *appears* to be informative about the world, but coincidentally leads to a true belief. For example, consider a scenario where a student comes to believe that his professor is in her office, because the student sees a convincing hologram sitting at the professor's desk. As it turns out, the professor is in her office, but she is crouching under the desk reading philosophy. In this case, the hologram serves as the evidence for the student's belief, which turns out to be true. However the hologram is not actually informative about whether the professor is in the room, and if the student realized that he was looking at a hologram, he would not have come to hold the belief that his professor was in her office. In contrast, the beliefs generated in the stories used in Experiments 1 and 2 are all based on authentic evidence: Peter forms the belief that

there is a watch on his table because he left a watch on his table; Jane believes that there is a \$20 bill in her cash register because she left one there. To investigate whether people consider Gettier cases based on authentic evidence to be knowledge, while viewing those based on apparent evidence to be mere belief, we next systematically manipulated the nature of the agent's evidence in Gettier cases.

6. Experiment 3

6.1. Method

6.1.1. Participants

Forty-three new participants (30 female, aged 19–82 yrs, mean = 36 yrs, standard deviation = 13 yrs) were tested. Data were excluded from an additional seven participants who failed control questions.

6.1.2. Materials and procedure

Participants read two stories, each from one of two conditions (Authentic Evidence or Apparent Evidence), and each presented using one of two storylines (coin or yogurt). In all stories, a character put an object in a location, and had a belief about the object's nature. For instance, in the “coin” stories, Corey puts a quarter dated 1936 in his piggy bank, and therefore believes his piggy bank contains a 1936 quarter. Although the coin is later removed by the his roommate, Corey's belief remains true because the piggy bank contains a different, previously unnoticed, 1936 quarter. What differed between conditions was whether Corey's initial belief was based on authentic evidence (i.e. the deposited coin really was a 1936 quarter) or apparent evidence (the deposited coin only looked like it was from 1936, but was actually a 1938 quarter).

Participants again responded to a comprehension question, a knowledge question, and a confidence question. If participants are sensitive to the authenticity of the evidence leading to the character's belief, knowledge attribution should differ between these two kinds of Gettier case. Across all participants, the storylines were presented equally in each condition, and order of presentation of storylines and conditions was counterbalanced. Story texts appear in the Appendix.




Fig. 3. Experiment 2. Left panel: Mean knowledge scores derived by combining dichotomous choice with confidence ratings. Right panel: Percentage of participants attributing knowledge in the Gettier and Low Justification conditions.

6.2. Results and discussion

A repeated measures ANOVA revealed a main effect of condition ($F(1,42) = 17.51, p < .001$), with participants attributing knowledge more often in the Authentic Evidence condition ($M = 3.70, SD = 8.03$) than in the Apparent Evidence condition ($M = -3.35, SD = 8.13, t(42) = -4.18, p < .001$, see Fig. 4). Knowledge attribution exceeded chance in the Authentic Evidence condition ($t(42) = 3.02, p = .004$) and knowledge was attributed at less than chance rates in the Apparent Evidence condition ($t(42) = -2.70, p = .01$). A similar, but more pronounced effect of condition is revealed when considering only the first story; participants were even more likely to attribute knowledge in the Authentic Evidence condition ($M = 4.90, SD = 7.17$), and even less likely to attribute knowledge in the Apparent Evidence condition ($M = -6.88, SD = 6.13, t(41) = -5.71, p < .001$). Confidence levels did not differ across conditions ($M = 8.6/10, t(42) = -.25, p = .81$).

Examining responses to the dichotomous knowledge question revealed a similar pattern: as shown in Fig. 4, participants attributed knowledge more in the Authentic Evidence condition (67%) than in the Apparent Evidence condition (30%, McNemar, $p = .001$). Knowledge attribution exceeded chance in the Authentic Evidence condition (binomial, $p = .03$) and was attributed at less than chance rates in the Apparent Evidence condition (binomial, $p = .01$). As before, this difference is magnified when considering only the first story; 76% of participants attributed knowledge in the Authentic Evidence condition, while only 14% attributed knowledge in the Apparent Evidence condition (Fisher's, $p < .001$). These findings reveal a difference between two kinds of Gettier case—as in the previous experiments, participants readily attributed knowledge when the Gettiered individual formed a belief based on authentic evidence. However, when the individual's belief was based on apparent evidence participants attributed mere belief.

7. General discussion

We examined people's folk conception of knowledge. Participants attributed knowledge to agents holding justified true beliefs, and denied that knowledge was possessed when beliefs were false (Experiments 1A and 1B) or

unjustified (Experiment 2). Participants viewed justified true beliefs as knowledge even when agents were Gettiered. Participants only denied that Gettiered agents were knowledgeable when the agents' beliefs were formed on the basis of apparent evidence (Experiment 3). Together, these findings suggest that: (1) people's folk conception of knowledge fits, at least roughly, with the classical philosophical view of knowledge as justified true belief, and conflicts with the current philosophical conception in which Gettiered individuals are viewed as not possessing knowledge; and (2) people do not consider beliefs based on apparent evidence to be knowledge, and so although there is a class of Gettier cases in which people do not attribute knowledge, this is not due to the agent being Gettiered.

The findings from the Gettier conditions are particularly striking because they contradict a longstanding assumption that Gettier cases obviously do not constitute knowledge. In light of these surprising results, we examine several alternative interpretations of our findings, and conclude that it is unlikely that the patterns of knowledge attribution we report are due to participants being confused about the stories or the questions, nor do they reflect peculiarities of our particular stories or methods.

One possible explanation is that participants misunderstood the test question, which asked whether the agent in each scenario “really knows” or “only believes” a certain proposition. For example, people might have interpreted “really knows” to mean “is very sure” or “thinks she knows”. Had participants interpreted the question this way, they should also have attributed knowledge to agents with false beliefs (because these agents had the same basis for being certain as the agents in the other conditions). But participants did not respond this way—instead they consistently denied that agents with false beliefs are knowledgeable.

A second possibility is that the unusual events described in the stories were confusing or unfamiliar to participants. This is unlikely; although the events described in these stories might seem unusual, Williamson (2005) points out that one can easily find examples of Gettier situations in real life. Moreover, even if participants did find the stories unusual, stories were closely matched across the different conditions; if unusual plotlines somehow led participants to attribute knowledge in the Gettier scenarios, high rates of knowledge attribution would likewise be expected across all conditions. However, this was not the pattern of results




Fig. 4. Experiment 3. Left panel: Mean knowledge scores derived by combining dichotomous choice with confidence ratings. Right panel: Percentage of participants attributing knowledge in the Authentic Evidence and Apparent Evidence conditions.

we observed. Also, if participants had been confused in the Gettier cases, they should have given low confidence ratings to their responses, but they did not. Confidence ratings did not differ across conditions, and moreover few participants ever used the lower end of the confidence scale.

Finally, perhaps peculiarities of the particular stories used were responsible for the differences in intuitions between laypeople and philosophers. However, participants in our studies showed similar patterns of knowledge attribution across multiple different storylines, and so the findings appear to be robust across a variety of contexts. Furthermore, it would be difficult to argue that participants simply did not notice that, in the Gettier scenarios, the fact causing the agents' belief to be true was causally disconnected from the evidence justifying the belief. Participants claimed that Gettiered agents were knowledgeable even after they answered a series of comprehension questions that highlighted this causal disconnect (Experiments 1A and 1B). Our findings suggest, then, people conceive of justified true beliefs as knowledge, and attribute knowledge to agents holding such beliefs even when the agent is Gettiered.

7.1. Apparent evidence

The conclusion that people view knowledge as justified true belief is complicated, though, by participants' reluctance to attribute knowledge when beliefs are based on apparent evidence—evidence that *appears* to be informative about reality, but is not really (Experiment 3). Consider the scenario in which Corey believed his piggy bank contained a quarter from 1936 because he deposited a quarter that *appeared* to be from 1936 (but was really a damaged quarter from 1938). Although Corey's belief was true (the piggy bank really did contain a quarter from 1936), and justified (seeing the date on a coin is usually appropriate justification for believing the coin is from that year), participants were less likely to attribute knowledge to Corey in this scenario than they were in a nearly identical scenario in which the quarter really was from 1936. This finding suggests that people will only view a belief as knowledge if it depends on evidence that is authentic rather than merely apparent. So rather than holding three conditions for knowledge (i.e., belief; justified; true), people may also hold a fourth “authenticity of evidence” condition.

This distinction between authentic and apparent evidence is largely intuitive; we have not described criteria for deciding when evidence is informative about reality (authentic evidence) and when it only appears to be informative (apparent evidence). A more precise way of drawing the distinction is to note that in cases of apparent evidence, the agent is unaware of facts about the evidence that if known would prevent the agent from forming the belief. For example, if Corey knew that the date on the quarter was damaged, he would not come to believe that the quarter was from 1936. An agent could likewise come to believe that it is 2 pm by seeing that it is 2 pm on a clock. But the agent would not come to believe this if the agent knew that the clock was broken, even if it really happened to be 2 pm. For undamaged coins and clocks there is no such unknown information which would prevent the belief from being formed. However, even this account is imprecise, and

counterexamples are probably possible; an important goal for future research and theorizing will be to come to a more formal account of when evidence is authentic or only apparent.

Regardless, this distinction is important because many of the Gettier cases discussed by philosophers involve apparent evidence (e.g., see the many cases reviewed in Lycan (2006)). Our results suggest that although laypeople might decline to attribute knowledge in these cases, it is only because these cases confound being Gettiered with the agent's belief being based on apparent evidence. This is exemplified by the findings of Experiment 4: people declined to attribute knowledge in Gettier scenarios featuring apparent evidence, but did attribute knowledge in Gettier scenarios featuring authentic evidence. Hence, the current research may point to a new factor that influences people's judgments about whether an agent has knowledge.

However, it could be that the justification and authenticity of evidence conditions on knowledge are not really distinct. This is not to claim that justification cannot be distinguished from authenticity of evidence. Rather the claim is that both may impact people's knowledge attributions for the same reason: People may adhere to a “no-luck” principle, holding that beliefs that are only true as a matter of luck do not count as knowledge, and this single principle might underlie the justification and authentic evidence conditions on knowledge. To see how this could be, first consider justification: Suppose that a friend glances at Corey's piggy bank, has a sudden hunch that it contains a quarter from 1936, and becomes convinced this must be true. The friend's belief is unjustified, and it is true only as a matter of luck. And this luck might be responsible for the intuition that the friend lacks knowledge. Now consider apparent evidence: When the quarter that Corey puts in the piggy bank only appears to be from 1936, (but is actually a damaged quarter from 1938), whether Corey's belief is true depends on luck—it depends on whether there are any other quarters from 1936 in the piggy bank. So again, people's judgment that Corey lacks knowledge might arise because they recognize that the truth of his belief depends on luck. Hence, even though justification and authenticity of evidence *can* be distinguished, people may be sensitive to both factors for a single reason—because beliefs that are unjustified or based on apparent evidence are only true as a matter of luck.⁵

7.2. Laypeople and philosophers

Probably the most striking finding from the current experiments is that laypeople readily attribute knowledge in Gettier cases. Whether this *seems* surprising will probably depend on one's own intuitions about whether knowledge is possessed in Gettier cases. But this finding suggests that laypeople have different intuitions from philosophers, who apparently hold with near unanimity that Gettiered agents are not knowledgeable (Williamson, 2005, p. 4). This difference would not be surprising if knowledge could

⁵ However, if people adhere to a no-luck principle, they should also be expected to view Gettiered individuals as lacking knowledge. We follow up on this implication below.

be observed, measured, or manipulated. Then it would be intuitive that those who study knowledge should come to have a special ability to identify it, just as it is intuitive that relative to laypeople, lepidopterists are better able to distinguish butterflies from moths, and radiologists are better able to interpret X-rays. But knowledge is not observable, except to the extent that people can reflect on their intuitions about what they believe counts as knowledge. So it is difficult to see how philosophers could have a special ability to identify instances of knowledge (see Weinberg, Gonnerman, Buckner, and Alexander (2010) and Williamson (2011) for more extensive debate about philosophers as experts). Why should laypeople and philosophers differ in their judgments about what counts as knowledge? We consider four explanations, which also imply different psychological accounts regarding how laypeople attribute knowledge.

One explanation is that the difference between laypeople and philosophers is akin to a novice–expert difference, even though knowledge is not a physical or external thing. As discussed above, people may adhere to the no-luck principle, which holds that a belief that is only true in virtue of luck cannot count as knowledge.⁶ If people adhere to this principle they should not attribute knowledge in Gettier cases, because it is a matter of luck that a Gettiered agent's belief is true. For example, when a thief takes Peter's watch and replaces it with another watch, it is a matter of luck that Peter's belief remains true—the thief could just as easily have replaced the watch with some other object, or Peter could just as easily have been robbed by a regular thief—one who does not replace stolen items with anything! So Peter's belief is only true by luck. The fact that people readily attribute knowledge to Peter might therefore seem to contradict the claim that they attribute knowledge according to the no-luck principle. However, it could be that people *do* adhere to this principle, but fail to apply it in Gettier cases because they do not recognize that Gettiered agents are only correct by luck. This failure could occur if laypeople typically assess for epistemic luck by considering the events *causing* the agent's belief to be formed, but without also considering that *subsequent* events that occur after belief formation (e.g., the thief replacing the watch) can also affect luck. Assessing for luck in this way would block knowledge from being attributed for true beliefs based on lucky guesses and apparent evidence, but not for true beliefs arising from Gettiering.

In contrast to laypeople, philosophers might be very practiced in thinking about the myriad ways in which the truth of a belief might rest on luck. Hence, laypeople and philosophers might both adhere to a no-luck principle, but nonetheless come to different conclusions about whether knowledge is possessed in Gettier cases. In this view, laypeople and philosophers might differ in performance, rather than competence, and this performance

difference might arise because laypeople are novices in detecting epistemic luck, while philosophers are experts; expertise does not require that knowledge is a physical or external thing. One reason to be skeptical about this explanation, though, is that participants attributed knowledge in Gettier cases even when asked very detailed questions about the scenario events (Experiments 1A and 1B). It is difficult to believe that participants answering these questions would fail to notice that the agent's belief was only correct by luck, because the questions highlighted that the justification for the belief was causally disconnected from the events making the belief true; moreover, the luck element was especially obvious in the 2-Thief Gettier case of Experiment 1B, where the second thief accidentally left the “replacement” pen on the table.

A second explanation for why laypeople do not share philosophers' intuitions is that laypeople and philosophers have different conceptions of knowledge. These two groups can be viewed as distinct cultural groups with distinct conventions regarding what counts as knowledge. On this view, the lay concept of knowledge involves holding a belief which is true, justified, and based on authentic evidence; however, training in philosophy might *enculturate* people into a new view of knowledge and induce a conceptual change. This explanation offers a very straightforward account of the present results, because it does not require a contrast between laypeople's performance (knowledge attributions to Gettiered agents) and their competence (adhering to a no-luck principle). Nor does it even require that people adhere to the no-luck principle. However, this account does not explain why philosophers happen to have a conception of knowledge that differs from the lay conception. Also, this account is challenged by the fact that philosophers' view of knowledge probably *preceded* the release of Gettier's paper. When the paper became public, it is claimed that most philosophers immediately agreed that Gettiered individuals do not possess knowledge (Williamson, 2005, p. 4). This suggests that rather than changing philosophers' intuitions about knowledge, Gettier's paper made evident the inability of the justified true belief theory of knowledge to account for philosophers' intuitions. However, if what counts as knowledge were a matter of convention, philosophers should have defended the justified true belief theory (rather than abandoning it), because it was the espoused view of their “culture”.

A third (related) explanation is that the difference between laypeople and philosophers arises because of selection effects. This view assumes that there are individual differences in laypeople's epistemic intuitions, and that some share the standard philosophical intuition while others do not. Those that have the “right” intuitions might be more likely to succeed in becoming philosophers; put another way, people who do not share the Gettier intuition might be unlikely to succeed as philosophers (or epistemologists, more specifically), and so the conception of knowledge shared by these laypeople should therefore be infrequent among philosophers (see Buckwalter & Stich, *in press*; Cummins, 1998, pp. 116–117). Like the enculturation view, though, this selection account does not explain *why* some people, and not others, view Gettiered individuals as lacking knowledge.

⁶ This discussion assumes that a “no luck” principle explains why philosophers claim that knowledge is not possessed in Gettier cases. Although “no luck” accounts have been prominent in attempts to explain why Gettiered individuals lack knowledge, they are not agreed upon by philosophers, and other proposals remain (for overviews of current accounts see Lycan (2006) and Turri (2012)). Even so, this experts and novices framework might also work for other principles claimed to explain why Gettiered agents lack knowledge.

A fourth explanation for why laypeople do not share philosophers' intuitions about Gettier cases is that the philosophical conception of knowledge might arise from epistemic theorizing. Suppose laypeople view knowledge as *justified true belief based on authentic evidence*, and that there is no deeper or more principled answer about what knowledge is. Philosophers, however, aim to provide an account for what knowledge *really is*, and in doing so they assume that there is a principled basis to knowledge—that it has a nature which can be discovered. The activity of trying to understand this principled basis may change how philosophers conceive of knowledge, and therefore also change how they attribute it. For example, insisting that knowledge requires absolute certainty, as Descartes did, may lead one to the conclusion that perceptual or sensory evidence is not sufficient for knowledge. Likewise, carefully considering why knowledge requires justification could lead one to the principle that knowledge should never be based on luck, and this in turn leads to the judgment that knowledge is not possessed in Gettier cases. But this might occur even if one's original conception of knowledge did not include the no-luck principle. Again, the act of theorizing about knowledge may change how one conceives of knowledge, and how one attributes it.

Regardless of whether any of these explanations is correct, the current experiments advance our understanding of how lay people conceive of knowledge. The findings suggest that the lay concept of knowledge is roughly consistent with the traditional description of knowledge as justified true belief, but with the caveat that people also require that the belief be based on authentic, rather than apparent, evidence. And the findings also suggest that laypeople's intuitions about what counts as knowledge differ from the intuition claimed to dominate among philosophers—a finding that either points to a competence-performance distinction in laypeople's knowledge attributions, individual differences in lay people's intuitions, or else the possibility that training or practice in philosophy induces conceptual change.

Acknowledgments

We thank Stephen Stich, Stephen Butterfill, Josh Knobe, and Alia Martin for their helpful comments and feedback on this paper. This work was funded by a grant from the Natural Sciences and Engineering Research Council of Canada awarded to OF.

Appendix A. Experiment 4

A.1. Scenario 1 – Yogurt

A.1.1. Authentic evidence

Julie buys a container of yogurt at the local deli. Although, Julie is not aware of it, the yogurt in the container is exceptionally sweet—a mixup at the factory caused the yogurt to get a triple dose of sweetener. Julie comes home, puts it her fridge, and then goes into her bedroom. Julie's neighbor Sam has been spying on her. While she is in her bedroom, he picks the lock to her apartment, and enters.

He takes the yogurt container from the fridge, and replaces it with a sealed container of yogurt from his own fridge. Then he goes back into his own apartment with Julie's yogurt container. Julie has only been in the bedroom for a few minutes, and did not hear anything.

At the end of the story:

1. Is there a container of yogurt in Julie's fridge? [**Yes/No**]
2. Julie _____ that there is a container of yogurt in her fridge. [Really knows/Only believes]

A.1.2. Apparent evidence

Julie buys a container of yogurt at the local deli. Although, Julie is not aware of it, there is no yogurt in the container—a mixup at the factory caused the container to be filled with sour cream instead. Julie comes home, puts it her fridge, and then goes into her bedroom. Julie's neighbor Sam has been spying on her. While she is in her bedroom, he picks the lock to her apartment, and enters. He takes the yogurt container from the fridge, and replaces it with a sealed container of yogurt from his own fridge. Then he goes back into his own apartment with Julie's yogurt container. Julie has only been in the bedroom for a few minutes, and did not hear anything.

At the end of the story:

1. Is there a container of yogurt in Julie's fridge? [**Yes/No**]
2. Julie _____ that there is a container of yogurt in her fridge. [Really knows/Only believes]

A.2. Scenario 2 – Coin

A.2.1. 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.

At the end of the story:

1. Is there a coin from 1936 in Corey's piggy bank? [**Yes/No**]
2. Corey _____ that there is a coin from 1936 in his piggy bank. [Really knows/Only believes]

A.2.2. 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.

At the end of the story:

1. Is there a coin from 1936 in Corey's piggy bank? [Yes/No]
2. Corey _____ that there is a coin from 1936 in his piggy bank. [Really knows/Only believes]

References

- Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, *116*(4), 953–970.
- Birch, S. A. J., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, *18*, 382–386.
- Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*, 1018–1034.
- Buckwalter, W. (2012). Non-traditional factors in judgments about knowledge. *Philosophy Compass*, *7*, 278–289.
- Buckwalter, W., & Stich, S. (in press). Gender and philosophical intuition. In J. Knobe & S. Nichols (Eds.) *Experimental philosophy* (Vol. 2). Oxford: Oxford University Press.
- Corriveau, K. H., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and recent accuracy. *Developmental Science*, *12*, 426–437.
- Cullen, S. (2010). Survey-driven romanticism. *The Review of Philosophy and Psychology*, *1*, 275–296.
- Cummins, R. (1998). Reflections on reflective equilibrium. In M. DePaul & W. Ramsey (Eds.), *The psychology of intuition and its role in philosophical inquiry* (pp. 113–127). New York: Rowman & Littlefield.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*, 121–123.
- Goldman, A. (1967). A causal theory of knowing. *Journal of Philosophy*, *64*, 357–372.
- Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, *17*, 757–758.
- Koenig, M. A., Clement, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, *15*, 694–698.
- Leslie, A. M., Friedman, O., & German, T. P. (2004). Core mechanisms in 'theory of mind'. *Trends in Cognitive Sciences*, *8*, 528–533.
- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, *73*, 1073–1084.
- Lycan, W. G. (2006). On the Gettier problem problem. In S. Hetherington (Ed.), *Epistemology futures* (pp. 148–168). Oxford: Clarendon Press.
- Nichols, S., & Stich, S. (2003). *Mindreading: An integrated account of pretence, self-awareness, and understanding other minds*. Oxford: Oxford University Press.
- Pratt, C., & Bryant, P. (1990). Young children understand that looking leads to knowing (so long as they are looking into a single barrel). *Child Development*, *61*, 973–982.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development*, *72*, 1054–1070.
- Sodian, B. (1988). Children's attributions of knowledge to the listener in a referential communication task. *Child Development*, *59*, 378–385.
- Sodian, B., & Wimmer, H. (1987). Children's understanding of inference as a source of knowledge. *Child Development*, *58*, 424–433.
- Sturgeon, S. (1993). The Gettier problem. *Analysis*, *53*, 156–164.
- Taylor, M., Cartwright, B. S., & Bowden, T. (1991). Perspective taking and theory of mind: Do children predict interpretive diversity as a function of differences in observers' knowledge? *Child Development*, *62*, 1334–1351.
- Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, *11*, 1–11.
- Turri, J. (2012). Gettier's wake. In S. Hetherington (Ed.), *Epistemology: The key thinkers*. New York: Continuum.
- Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, *65*, 157–170 (reprinted in Roth and Galis (1970)).
- Weinberg, J. M., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology*, *23*, 331–355.
- Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, *29*, 429–460.
- Williams, M. (1978). Inference, justification, and the analysis of knowledge. *Journal of Philosophy*, *75*, 249–263.
- Williamson, T. (2005). Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings of the Aristotelian Society*, *105*, 1–23.
- Williamson, T. (2011). Philosophical expertise and the burden of proof. *Metaphilosophy*, *42*, 215–229.
- Wimmer, H., Hogrefe, G. J., & Perner, J. (1988). Children's understanding of informational access as source of knowledge. *Child Development*, *59*, 386–396.
- Woolley, J. D., & Wellman, H. M. (1993). Origin and truth: Young children's understanding of imaginary mental representations. *Child Development*, *64*, 1–17.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *The Philosophical Quarterly*, *44*, 65–73.