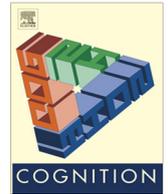




ELSEVIER

Contents lists available at SciVerse ScienceDirect

## Cognition

journal homepage: [www.elsevier.com/locate/COGNIT](http://www.elsevier.com/locate/COGNIT)

## Discussion

## Taking 'know' for an answer: A reply to Nagel, San Juan, and Mar

Christina Starmans<sup>a,\*</sup>, Ori Friedman<sup>b,\*</sup><sup>a</sup> Department of Psychology, Yale University, USA<sup>b</sup> Department of Psychology, University of Waterloo, Canada

## ARTICLE INFO

## Article history:

Available online xxxx

## Keywords:

Knowledge  
Gettier cases  
Lay theories  
Folk psychology  
Theory of mind  
Epistemic intuitions

## ABSTRACT

Nagel, San Juan, and Mar report an experiment investigating lay attributions of knowledge, belief, and justification. They suggest that, in keeping with the expectations of philosophers, but contra recent empirical findings [Starmans, C. & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124, 272–283], laypeople consistently deny knowledge in Gettier cases, regardless of whether the beliefs are based on 'apparent' or 'authentic' evidence. In this reply, we point out that Nagel et al. employed a questioning method that biased participants to deny knowledge. Moreover, careful examination of participants' responses reveals that they attributed knowledge in Gettier cases. We also note that Nagel et al. misconstrue the distinction between 'apparent' and 'authentic' evidence, and use scenarios that do not feature the structure that characterizes most Gettier cases. We conclude that NS&M's findings are fully compatible with the claim that laypeople attribute knowledge in Gettier cases in general, but are significantly less likely to attribute knowledge when a belief is generated based on apparent evidence.

© 2013 Elsevier B.V. All rights reserved.

## 1. Introduction

Which claims should be considered mere beliefs and which count as knowledge? For thousands of years, philosophers have sought to answer this question through debate. Although this debate continues, a general consensus has emerged among philosophers: Beliefs typically count as knowledge when they are justified and true, but justified true beliefs can fall short of knowledge in scenarios termed 'Gettier' cases. Recently, psychologists and experimental philosophers have begun running experiments to discover the conditions under which people attribute knowledge. In a recent paper taking this approach, Starmans and Friedman (2012; henceforth S&F) found a conflict between laypeople's intuitions about knowledge

and the views espoused by philosophers. Although laypeople share the intuition that justified true beliefs typically count as knowledge, they do not deny knowledge in Gettier cases. People do deny knowledge in a subset of Gettier cases, where the agent's belief is based on 'apparent' evidence rather than 'authentic' evidence. But this suggests that they are sensitive to the nature of evidence justifying beliefs, rather than 'Gettiering' more generally.

Nagel, San Juan, and Mar (this issue; henceforth NS&M) report an experiment investigating lay attributions of knowledge, belief, and justification. Their findings provide support for previous findings that laypeople recognize a difference between knowledge and justified true belief (see also Buckwalter & Stich, 2010; Starmans & Friedman, 2012; Turri, forthcoming; Weinberg, Nichols, & Stich, 2001). They also provide the first confirmation that, as has been only assumed in previous work, laypeople do in fact consider the beliefs described in these scenarios to be justified. Furthermore, they replicate recent findings that knowledge attributions are not generally affected by demographic variables (Starmans & Friedman, 2012; Turri, forthcoming; Wright, 2010; but see Buckwalter, 2010;

\* Corresponding authors. Addresses: Department of Psychology, Yale University, 2 Hillhouse Ave, New Haven, CT 06511, USA. Tel.: +1 203 432 0877 (C. Starmans), Department of Psychology, University of Waterloo, 200 University Ave West, Waterloo, Ontario, Canada N2L 3G1. Tel. +1 519 888 4567x33054 (O. Friedman).

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

Weinberg et al., 2001), though they find evidence for the interesting suggestion that more empathetic individuals are more likely to attribute knowledge to others. Finally, contrary to S&F, but in keeping with the expectations of philosophers, they claim that laypeople consistently deny knowledge in Gettier cases, including both cases where belief are based on ‘apparent’ or ‘authentic’ evidence.

In this reply we make three points relevant to understanding this last, most striking, finding. First, although NS&M suggest that laypeople deny knowledge in authentic evidence Gettier cases, the authentic evidence cases they tested feature a different structure than most Gettier cases, including the cases used by S&F. As a result, although these cases may be interesting to examine, knowledge attribution in these cases is not informative about how laypeople reason about most Gettier cases. Second, NS&M inaccurately equate the distinction between ‘apparent’ and ‘authentic’ evidence raised by S&F with the distinction between ‘false lemmas’ and ‘no false lemmas’, raised by Clark (1963; see also Harman, 1973). Finally, NS&M employed a questioning method that biased participants to deny knowledge. However, careful examination of their responses reveals that participants attributed knowledge at high rates in Gettier cases, and perhaps at higher rates than in S&F. We conclude that NS&M’s findings are fully compatible with the claim that laypeople attribute knowledge in Gettier cases in general, but are significantly less likely to attribute knowledge when a belief is generated based on apparent evidence.

## 2. The structure of Gettier cases

Why care about whether laypeople attribute knowledge in Gettier cases? Gettier’s initial goal in raising this type of case was as an existence proof against the classical view that any justified true belief (JTB) counts as knowledge. The vast majority of philosophers now accept this, and that laypeople also sometimes deny knowledge in cases of JTB is by now well established. In fact, a variety of factors beyond truth and justification have been shown to affect knowledge attributions, including stakes, pragmatic load, morality, performance errors, and demographic variation (e.g., Buckwalter, 2012 for review). And in their current paper, NS&M find that laypeople fail to attribute knowledge in a distinct set of cases in which there is no ‘Gettiering’, but merely a reminder that it might be prudent to be skeptical about knowledge.

What continues to be interesting about Gettier cases, then, is that they provide the opportunity to test which factors affect knowledge attributions. Gettier’s original two cases, and most all cases to follow, share a common structure. They describe a situation in which a person has a belief that is both justified and true, but the fact that justifies the belief is not the same fact that makes the belief true. For instance, S&F provide an example case in which Peter places his watch on his coffee table. While Peter is out of the room briefly, a burglar steals his watch and replaces it with another watch. Peter formed the belief that there was a watch on his table because he put *his* watch on the table, but at the end of the scenario, Peter’s belief is true

for a completely different reason—it is true because the burglar left a different watch on the table. So examining intuitions about Gettier cases allows us to test whether people’s knowledge attributions are sensitive to this disconnect between the fact that justifies the belief and the fact that makes it true.

However, NS&M base their authentic evidence cases on a class of cases that do not feature the disconnect structure that characterizes most Gettier cases. Their cases are modeled on Ginet’s “Fake Barn” cases (Goldman, 1976). The original case describes a rural landscape which appears to be populated with many barns. Looking at one of these objects, Henry believes he’s looking at a barn. And, in fact, Henry is looking at a barn. However the barn he is looking at is the only genuine barn in the whole area—the rest are very convincing facades, made to look like barns. In this case, the evidence for the agent’s belief (i.e., seeing a particular barn in the field) is the *same* fact that makes the agent’s belief true. So although such ‘fake barn’ cases are often referred to as Gettier cases, they do not feature the disconnect characteristic of most other Gettier cases. And in fact, philosophers themselves are quite divided on whether to attribute knowledge in these cases (e.g., Lycan, 2006; Sosa, 2007; Turri, 2012).

As a result, examining knowledge attribution in these cases does not shed light on whether people are sensitive to the factor common to most other Gettier cases—the disconnect between the fact that justifies the belief and the fact that makes the belief true. In contrast, S&F provide evidence that laypeople are not sensitive to this factor; in 3 studies participants attributed knowledge in scenarios featuring this disconnect. A fourth study confirmed that while people sometimes do deny knowledge in Gettier scenarios, it is not because of the disconnect described above, but because of the nature of the evidence justifying the original belief.

## 3. Evidence and lemmas

NS&M report that laypeople deny knowledge in two types of Gettier cases, which they refer to as ‘authentic evidence’ and ‘apparent evidence’ (adopted from S&F). They suggest that this distinction between apparent and authentic evidence is identical to the distinction between false lemmas and no false lemmas (Clark, 1963). More specifically, they suggest that apparent evidence cases are equivalent to false lemma cases, where an agent arrives at a true belief by reasoning on the basis of false steps, and that authentic evidence cases are equivalent to no false lemma cases, where no false steps occur in the agent’s reasoning. However, these two distinctions are not equivalent.

S&F introduce the terms authentic and apparent evidence to capture a distinction between two sorts of evidence an agent can have when initially *forming* a belief. Authentic evidence is informative about how the world actually is when the belief is formed, and basing a belief on authentic evidence necessarily makes the belief true when it is formed. Apparent evidence only appears to be informative about how the world actually is, and basing a belief on apparent evidence does not guarantee that the belief is true when it is formed.

For example, NS&M give us a case in which Luke believes that one of his co-workers, Victor, went to Las Vegas on his vacation because Victor has been deliberately misleading him with fake photos, etc. Victor did not actually go to Vegas, but another of Luke's co-workers, Monica, did. Luke's belief that one of his co-workers went to Vegas is justified by the extensive evidence provided by Victor. In this case, the evidence for Luke's belief is only apparent—the pictures he saw *seemed* to confirm that his co-worker had vacationed in Vegas, but basing a belief on this evidence does not ensure that the belief is true (when the belief is formed). And if Luke knew more about the nature of the evidence, he would not hold this belief. In contrast, in the case with Peter's watch discussed earlier, Peter forms his belief based on the fact that he had just put his watch on the table. Here, the evidence is authentic—the watch Peter is looking at really is a watch, it really is on the table, and following this evidence guarantees that his belief will be true (at least when he forms it).

Conversely, the distinction between 'false lemmas' and 'no false lemmas', as NS&M note, is whether the knowledge has arisen "from reasoning that essentially relies on a false belief". Thus apparent evidence cases necessarily contain false lemmas—the evidence gives rise to a false belief about the world (i.e., 'these photographs depict Victor in Las Vegas'), which can coincidentally give rise to a true belief (i.e., 'one of my coworkers went to Las Vegas'). Authentic evidence cases, however, may *also* contain false lemmas. The nature of the evidence used when an agent initially forms a belief says nothing about whether *maintaining* the belief depends on subsequent false beliefs or assumptions (i.e., false lemmas). In fact, all the authentic evidence Gettier cases used in S&F feature false lemmas. For example, in the case of Peter's watch, Peter initially forms the true belief that a watch is on the table from authentic evidence. However, by the end of the story, his belief depends on the false assumption that *his* watch is where he left it.

NS&M claim to provide evidence that, contra S&F, laypeople attribute knowledge in Gettier cases in which the agent's belief is based on authentic evidence. However none of their 8 cases fall into this category. Four cases are correctly described as apparent evidence cases. The remaining four cases are classified as containing both authentic evidence and no false lemmas. However, three of the four 'authentic' cases NS&M tested were actually apparent evidence cases. For example, one scenario is a 'stopped clock' case (Russell, 1948; see S&F Section 7.1 for discussion), in which Wanda looks at a clock and forms a belief that it is 4:15. It is in fact 4:15, but the clock has been stopped for the past 2 days, and just happens to show the correct time. This is a clear case of apparent evidence: Heeding a broken clock does not guarantee that her belief is true when it is formed, and if Wanda knew the true nature of the evidence she was relying on, she would not come to believe that it is 4:15.<sup>1</sup> Thus, NS&M do not provide evidence that laypeople deny knowledge in authentic evidence Gettier cases.

<sup>1</sup> Additionally, it seems that most of the 'no false lemma' cases NS&M tested, like this one, also contain false lemmas. That is, Wanda's belief that it is 4:15 essentially relies on the false belief that the clock is in working order.

#### 4. Taking "know" for an answer, and other methodological concerns

Our third concern centers on the method NS&M used to question participants about knowledge. After participants read each vignette, they were initially asked whether the protagonist was knowledgeable. If participants denied knowledge, their response was accepted. However, if participants attributed knowledge to the protagonist, they were asked a follow-up question, asking them to choose between whether the protagonist *knows* or only *feels like she knows, but does not actually know*. This questioning method allowed participants to change their responses if they attributed knowledge, but not if they denied knowledge, increasing the likelihood of knowledge denial. In fact, if all participants responded at chance, then this procedure would result in only 25% of participants unwaveringly attributing knowledge.<sup>2</sup> However, NS&M found unwavering knowledge attributions in 35.4% of responses (see their Table 2). So although they claim that their findings suggest that people deny knowledge in Gettier cases, a closer examination reveals that people attributed knowledge at rates *exceeding* chance levels (Binomial test with chance at 25%,  $N = 263$  responses,  $k = 93$  unwavering knowledge attributions,  $p < .001$ , two-tailed). These high rates of knowledge attribution are somewhat surprising, given that almost all of their cases featured apparent evidence in some form. If anything, NS&M probably found *higher* rates of knowledge attribution in these cases than did S&F! This high rate of knowledge attribution is also surprising because repeatedly being asked follow-up questions might have led participants to assume that their initial answers were incorrect or unsatisfactory, and might have pressured them to deny knowledge.<sup>3</sup>

NS&M explain that they used follow-up questions because participants could have interpreted the initial question about knowledge in a "projected or non-literal sense", taking the question to concern whether the protagonist *thinks* that she knows, even if she literally does not. The concern that participants could have misinterpreted the question, *on first encountering it*, is legitimate. However, other steps could have been taken to avoid the prospect of participants misinterpreting the meaning of "knows", while simultaneously avoiding pressuring subjects to change their answer only if they attributed knowledge.<sup>4</sup> Regardless, it is very unlikely that participants could have continued to misinterpret the initial knowledge questions

<sup>2</sup> For example, in a group of 100 participants responding at chance, 50 would initially deny knowledge and their responses would stand. The other 50 would initially attribute knowledge, but when offered the chance to deny knowledge 25 would switch, leaving only 25 unwaveringly attributing knowledge.

<sup>3</sup> NS&M discuss several methodological limitations of their study in Section 3, and some of these might explain the high rates of knowledge attribution. These limitations include the complexity of the scenarios, ambiguity about *when* knowledge is being attributed, and the fact that participants responded to 16 vignettes. Previous studies (Wright, 2010) have found that knowledge attributions vary wildly depending on the order in which participants had read the stories—as much as 20% greater knowledge attribution for the same story.

<sup>4</sup> For example, participants could have been asked only the follow-up questions, without the initial questions asked at all, because these questions make clear the sense in which 'know' is used.

after receiving their first follow-up question. Once the meaning of “know” was clarified by a follow-up question, it is difficult to see how participants could continue to misinterpret the initial knowledge questions in the subsequent vignettes.

In fact, NS&M’s results strongly suggest that, on the whole, participants did not misinterpret the initial knowledge questions. If participants had interpreted the questions as asking about whether the protagonist *feels* that she knows (rather than as concerning whether she *actually* knows), then participants should have initially attributed knowledge in the false belief vignettes. In these vignettes, the protagonists lacked knowledge, but had equal reason to feel they possessed knowledge as in the other vignettes. Contrary to this prediction, participants denied knowledge in the false belief vignettes, even when asked the initial knowledge questions—only 32.2% of initial responses were knowledge attributions.

Given this evidence that participants did not misinterpret the initial knowledge questions, their responses to these questions are the best place to look for participants’ *actual* knowledge attributions; unlike the follow-up questions, these questions did not pressure participants to deny knowledge. And examining these questions, we again see that participants *did* attribute knowledge in the Gettier cases (59.9% knowledge attribution; Binomial test with chance at 50%,  $N = 263$  responses,  $k = 108$ , knowledge denials,  $p = .002$ , two-tailed). As noted above, these high rates of knowledge attribution are surprising given that most of the cases featured apparent evidence. Rather than showing that laypeople’s knowledge attributions are in line with philosophical intuition, NS&M’s findings suggest that people attribute knowledge in Gettier cases. However, even

this conclusion should be interpreted with caution due to the concerns discussed above.

### Acknowledgment

We thank John Turri for helpful comments and feedback on this paper.

### References

- Buckwalter, W., & Stich, S. (2010). Gender and philosophical intuition. SSRN. <http://ssrn.com/abstract=1683066>. Revised version to appear in Knobe, J., & Nichols, S. (Eds.), *Experimental Philosophy*, vol. 2. Oxford: Oxford University Press.
- Buckwalter, W. (2012). Non-traditional factors in judgments about knowledge. *Philosophy Compass*, 7, 278–289.
- Clark, M. (1963). Knowledge and grounds: A comment on Mr. Gettier’s Paper. *Analysis*, 24(2), 46.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Lycan, W. G. (2006). On the Gettier problem problem. In Stephen Hetherington (Ed.), *Epistemology Futures* (pp. 148–168). Oxford: Clarendon Press.
- Nagel, J. San Juan, V. & Mar, R. (2013). Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, xxx–xxx.
- Russell, B. (1948). *Human knowledge: Its scope and its limits*. New York: Simon & Schuster.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford: Oxford University Press.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124, 272–283.
- Turri, J. (2012). Is knowledge JTB? *Synthese*, 184(3), 247–259.
- Turri, J. (forthcoming). A Conspicuous Art: Putting Gettier to the Test, *Philosophers’ Imprint*.
- Weinberg, J. S., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1), 429–460.
- Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115(3), 491–503.