Adrian Ziółkowski

University of Warsaw

adrian.a.ziolkowski@uw.edu.pl

[draft version, please do not cite]

[forthcoming in *Oxford Studies in Experimental Philosophy* vol. 4]

Folk Knowledge Attributions and the Protagonist Projection Hypothesis[1]

ABSTRACT

A growing body of empirical evidence suggests that folk knowledge attribution practices regarding some epistemological thought experiments differ significantly from the consensus found in the philosophical literature. More specifically, laypersons are likely to ascribe knowledge in the so-called Authentic Evidence Gettier-style cases, while most philosophers deny knowledge in these cases. The intuitions shared by philosophers are often used as evidence in favor (or against) certain philosophical analyses of the notion of knowledge. However, the fact that these intuitions are not universal, as non-philosophers do not sympathize with verdicts made by philosophers, is problematic and requires some explanation. Recently, a promising theoretical approach emerged that could be used to explain away unexpected folk knowledge attributions. According to the protagonist projection hypothesis, when subjects answer questions about some hypothetical scenario, their judgments might result from adopting the cognitive perspective of the protagonist in the given scenario. Previous experimental findings suggest protagonist projection might be responsible for problematic knowledge attributions in Gettier-style cases. This paper reports data collected in an experiment which focused on five different Authentic

Evidence Gettier-style scenarios and investigated the impact of protagonist projection on folk knowledge ascriptions in such cases. The results bring some support to the protagonist projection hypothesis, but do not allow to explain away a substantial number of problematic folk knowledge attributions.

1. Introduction

Many epistemologists believe that counterexamples to the classic theory of knowledge proposed by Edmund Gettier (1963) provide an ultimate refutation of that theory. According to the classic (or tripartite) account of knowledge, knowledge possession involves three simple constituents: 1) belief, 2) truth, and 3) justification. In terms of necessary and jointly sufficient conditions: for any proposition *p*, and any agent *A*, *A* knows that *p* iff *A* has a true, justified belief that *p*. Contrary to that idea, Gettier cases describe situations in which agents possess certain true and justified beliefs—which, according to the classic analysis, should count as knowledge—but nevertheless elicit intuitions that the agents in question do not know the relevant propositions. Since then Gettier scenarios (and other similar cases introduced later that I will call *Gettier-style cases*, since they also aim at eliciting intuitions inconsistent with the classic theory of knowledge) have become the crucial test that any proposed analysis of knowledge needs to pass in order to receive recognition.

However, recently some researchers (e.g., Starmans & Friedman 2012) argued that it is unclear that Gettier (and Gettier-style) cases unequivocally elicit intuitions shared by epistemologists. It turned out that many subjects not trained in philosophy tend to attribute knowledge in some situations of that sort. While this may not be the case for scenarios modeled on the original Gettier cases, where most non-philosophers deny knowledge in accordance with philosophers (e.g. Nagel et al. 2013a; Machery et al. 2017a, 2017b), it is a serious issue for other, more elaborate Gettier-style cases (similar to the *Fake Barns* scenario introduced by Goldman, 1976), where the majority of subjects are happy to attribute knowledge against philosophical consensus (e.g. Colaço et al. 2014; Ziółkowski 2016). The agreement between philosophers and non-philosophers regarding the former type of cases (sometimes called *Apparent Evidence*

cases) is believed to be unsurprising, since here the crucial belief is a result of an impaired cognitive process (e.g., valid reasoning from false premises), which may violate the justification requirement. The latter type, also known as *Authentic Evidence cases*, does not have that trait.

One interesting attempt to tackle these troublesome knowledge attributions rests on the reference to a phenomenon known as protagonist projection, which consists in a tendency to answer questions from the cognitive perspective of the protagonist in the presented scenario (for a detailed definition of this phenomenon, see Holton, 1997). Some tentative data suggest that protagonist projection might in fact play an important role here—many subjects who attribute knowledge in Gettier cases do so because they adopt the perspective of the protagonist in question (Nagel et al. 2013a; Machery et. al 2017a, 2017b). However, most of the data collected so far concern Apparent Evidence cases.

I will focus on Authentic Evidence cases and aim at establishing whether the strong tendency of laypersons to attribute knowledge in such cases stems from subjects' inclination to project into the cognitive perspective of the protagonists in the scenarios. If that turned out true, folk judgments concerning these cases could not be treated as genuine knowledge attributions and would not be inconsistent with philosophical tradition.

In section 2, I briefly sketch previous empirical findings regarding folk knowledge attribution practices in Gettier-style cases, with special attention devoted to results that seem to be at odds with the consensus found in the philosophical literature. In section 3, I introduce the protagonist projection hypothesis and discuss its possible explanatory merits with regards to unexpected folk knowledge ascriptions. Then, in section 4, I turn to present the methods employed in my experiment and the results it yielded. My findings are rather unfavorable to the protagonist projection hypothesis—while there is some level of projection in subjects' answers,

for some Authentic Evidence cases I investigated the majority of participants tend to attribute knowledge even if they are offered an opportunity to withdraw their antecedent attributions and indicate protagonist projection. Therefore, in section 5, I conclude that it seems that a substantial number of problematic folk knowledge ascriptions cannot be explained away by protagonist projection.

2. Troublesome experimental findings

The original scenarios presented by Gettier (1963) describe situations in which an agent forms a belief (say: *B*) that is true and justified, but it seems wrong to say that *B* is an instance of knowledge. However, the belief in question always results from a deductive reasoning based on a false premise. For example: Smith has strong evidence that Jones owns a Ford and forms an according belief. Then, Smith performs a perfectly valid reasoning based on disjunction introduction rule—if *p*, then *p*∨*q*—and forms another belief that Jones owns a Ford or Brown is in Barcelona. This latter belief turns out to be true, but thanks to the fact that Brown is in Barcelona—Smith happens to be wrong about Jones (he actually drives a rental), but luckily introduced a true disjunct about Brown. Thus, Smith's true and justified belief that either Jones owns a Ford or Brown is in Barcelona resulted from a valid reasoning based on a false premise (that Jones owns a Ford). Starmans & Friedman (2012) propose to call this type of Gettier-style cases *Apparent Evidence cases*, since the crucial belief in this kind of scenarios is always based on apparent (false) evidence and the process of belief acquisition is cognitively flawed. These kinds of scenarios are also referred to as *false lemma cases*, since they involve an inference from a false premise, which makes the reasoning in question obviously incorrect (although it is correct formally, it is materially incorrect). For this reason, it would not be surprising to discover that

subjects not trained in philosophy deny knowledge in such cases, similarly to professional philosophers.

Empirical evidence suggests that Gettier's own intuitions about Apparent Evidence cases are in fact shared by many laypersons. For example, Machery et al. (2017a) in their cross-cultural study that investigated two different Apparent Evidence cases, found that folk intuitions elicited by these cases are in line with those expressed in philosophical literature for most investigated ethnic groups (participants from India were an exception here). For example, for their *Hospital* case, Machery et al. (2017a) found that the majority of respondents from USA (70.3%) denied knowledge straight away, when offered a simple "Yes-or-No" knowledge question. The results collected from Brazilian and Japanese subjects were similar. A follow-up study conducted by Machery et al. (2017b) that investigated users of 17 different languages corroborated the initial findings—Gettier intuitions seem to be widespread and quite consistent across ethnicities.

Unfortunately for philosophers, Apparent Evidence cases are not the only type of Gettier-style cases important for epistemological debates. There is another class of Gettier-style cases, sometimes called *Authentic Evidence cases* (e.g. Starmans & Friedman, 2012), that describe situations where an agent has a true, justified belief that results from solid (true) evidence.[2] Nevertheless, most philosophers would deny knowledge in such cases. One of the most famous examples of that sort is the *Fake Barns* case authored by Carl Ginet and first introduced in the literature by Alvin Goldman (1976). One variant of the scenario used in the experiment conducted by Colaço et al. (2014) is as follows:

Gerald is driving through the countryside with his young son Andrew. Along the way he sees numerous objects and points them out to his son. "That's a cow, Andrew," Gerald says, "and that over there is a house where farmers live". Gerald has no doubt about what the objects are. What Gerald and Andrew do not realize is the area they are driving through was recently hit by a very serious tornado. This tornado did not harm any of the animals, but did destroy most buildings. In an effort to maintain the rural area's tourist industry, local townspeople built house façades in the place of destroyed houses. These façades look exactly like real houses from the road, but are only for looks and cannot be used as actual housing. Though he has only recently entered the tornado-ravaged area, Gerald has already encountered a large number of house façades. However, when he tells Andrew "That's a house", the object he sees and points at is a real house that has survived the tornado (Colaço et al. 2014, p. 201).

In this scenario, Gerald has a true and justified belief that he saw a house, but it is true only due to luck—he could have easily looked at a house façade and formed a false belief instead.[3] In the Gettierized condition the researchers asked their respondents "Does Gerald know he saw a house?", while in the control condition the question was "Does Gerald know he saw a cow?". Answers were measured on a six-point Likert scale. Colaço et al. (2014) observed a significant difference between the control and Gettierized conditions (according to benchmarks proposed by Cohen, 1988, the size of the observed effect, $d = 0.39$, could be classified as small-to-medium); however the average knowledge rating in the Gettierized condition was still significantly above the midpoint of the scale. Thus, contrary to the widespread consensus among

philosophers, most of laypersons were likely to attribute knowledge in the *Fake Barns* case. These findings were later corroborated by Ziółkowski (2016).

Another important Gettier-style thought experiment discussed in the literature, which I will refer to as the *Thermometer* case, was put forward by Goldman (1986). It was first adopted for the sake of empirical research and experimentally tested by Ziółkowski (2016). The version he used is presented below.

Stephen is a high school history teacher. One morning, Josh, his son, enters the kitchen. Josh feels ill. He complains about having a cough, runny nose, and a rash on his skin. Stephen examines the symptoms of Josh's disease in detail. He discovers that most of his son's body is covered with small red spots. He also sees that Josh has red, bloodshot eyes. Stephen wants to check whether Josh also has a fever. Stephen takes a box full of thermometers out of the kitchen cupboard. He picks one thermometer from the box and takes Josh's temperature. The thermometer indicates 98.6 degrees, leading Stephen to believe that his son's temperature is normal, which is true. However, unbeknownst to Stephen, the thermometer he used is the only reliable thermometer in the box. All other thermometers in the box are defective and they would read 98.6 even if Josh had a fever. Still, the thermometer that Stephen used was working fine and gave the correct reading of the temperature (Ziółkowski, 2016, p. 357).

Although the comparison between the control condition ("Does Stephen know his son has red eyes?") and the Gettierized condition ("Does Stephen know his son's temperature is 98.6?")

yielded statistically significant results, the size of the difference was not very large ($d = 0.49$). Most laypersons judged that Stephen knows his son's temperature, regardless of the fact pointed out by many philosophers, that here the agent was lucky to pick the only reliable thermometer from the box. The presence of epistemic luck did not stop laypersons from attributing knowledge in this case.

Starmans & Friedman (2012) tested yet another kind of Authentic Evidence cases, where the belief in question is true and well-justified, but its truthmaker is being replaced in the process. Here is one of their vignettes, which I will call *Wristwatch*:

> 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. 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, and then leaves. Peter is still in the shower, and did not hear anything (Starmans & Friedman, 2012, p. 274).

The respondents were asked to fill in the blank space in the statement "Peter ... that there is a watch on the table" (or "Peter … that there is a book on the table" in the control condition) by choosing from two options: "really knows" or "only thinks". Once again, the majority of laypersons did not find a good reason to deny knowledge in the Gettierized condition, where 72% of subjects attributed knowledge (compared to 88% in the control condition).[4] Starmans &

Friedman (2012) obtained similar results for another Authentic Evidence case they tested in their second experiment (the *Bic Pen* scenario, which can be found in Appendix A).

All the empirical findings discussed above may pose serious problems for epistemologists. Of course the relationship between folk knowledge attribution practices and philosophical theories of knowledge may not be a straightforward one—in particular, the discovery that some philosophical analyses of knowledge are not backed up by folk intuitions does not have to mean that these analyses should be discarded. Nevertheless, the fact that, contrary to a widespread consensus among professional philosophers, laypersons tend to attribute knowledge in Authentic Evidence Gettier-style cases, requires some explanation. Starmans & Friedman (2012) put forward a bold conclusion that the folk conception of knowledge is simply different from the notion of knowledge shared by epistemologists. However, some philosophers consider this conclusion unwelcome, as they believe that the notion of knowledge should rather be universal. One way to preserve the universality thesis is to provide an error theory of these folk epistemic intuitions which are at odds with philosophical tradition. Such a theory, pointing at possible biases in laypersons' judgments or misunderstandings these judgments stem from, might be used to explain away troublesome data presented above, thus, making them harmless for the universality claim. One theoretical approach following this line of reasoning, the protagonist projection hypothesis, is discussed in the next section.

3. The Protagonist Projection Hypothesis

The protagonist projection hypothesis is a promising theoretical account that might be helpful in dealing with the puzzling empirical findings presented above. According to this hypothesis, when we confront subjects with a scenario describing some agent and ask them questions

regarding the scenario, the answers provided by the subjects might result from their tendency to adopt the cognitive perspective of the protagonist in the scenario. To put it differently, respondents may project the cognitive perspective of the protagonist onto their own answers, thus, the name *protagonist projection hypothesis*. It was first presented by Holton (1997) and later used in experimental philosophy (among others) by Buckwalter (2014) in his research concerning seemingly non-factive uses of factive verbs (such as 'know', 'learn', or 'realize'). I will now present some of Buckwalter's findings to illustrate the possible explanatory merits of the protagonist projection hypothesis.

Factive verbs are usually used in language expressions to refer to dyadic relations occurring between agents and propositions—for instance, a sentence '*A* knows that *p*' refers to a knowledge relation holding between the agent *A* and the proposition *p*. Traditionally, it is assumed that statements containing such verbs logically entail other factual claims—for example, if '*A* knows that *p*' is true, then the proposition *p* has to be true as well (similarly for 'realize' or 'learn')—and for that reason we call them *factive* verbs. However, in ordinary language one can easily find examples of statements which seem to violate the factivity assumption. Here is one example used in Buckwalter's (2014) experiments: (A) "Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection". While (A) sounds like a potentially felicitous claim, it also suggests that it is possible to know a proposition *p* even when *p* is false, so it appears to be a counterexample to the factivity assumption. However, Buckwalter presents an alternative interpretation: it might be that when subjects make (or accept) utterances like (A), they attach some non-literal meaning to them that stems from projecting into someone else's cognitive perspective. In the case of (A), subjects might in fact mean that before it had been proven that

ulcers is caused by bacterial infection, people *thought they knew* stress caused ulcers, but they were simply wrong. Buckwalter presented his respondents with sentences like (A), and asked them what is meant by 'knew' in the sentence, offering them two options: (a) "Everyone thought they knew"; (b) "Everyone really did know".[5] It turned out that only 9% indicated (b), while the rest chose the projective answer, which suggests that they adopt a non-literal reading of (A).[6] This non-literal meaning is not inconsistent with the factivity of 'know'. Therefore, thanks to the protagonist projection hypothesis, we can show that claims such as (A) are only apparent counterexamples to the factivity assumption.

A similar argumentative strategy to that presented by Buckwalter (2014) in the discussion about the factivity assumption might be formulated in order to explain away problematic knowledge attributions, which is the main issue investigated in this paper. After all, it is possible that knowledge ascriptions made by laypersons in Authentic Evidence Gettier-style cases are only apparent and result from protagonist projection. In other words, when subjects judge that the protagonist of a given scenario knows that p, they might really mean that the protagonist thinks he or she knows p, but in fact does not know it. If that was the case, high ratios of knowledge attributions made by laypersons in Authentic Evidence cases (see section 2) would not be problematic for epistemology, as they would not be genuine knowledge attributions in the first place. This idea was put to the test by Nagel et al. (2013a), who examined a number of both Apparent and Authentic Evidence cases. Participants of their study were given two opportunities to deny knowledge—first in a straightforward question ("Does the protagonist know whether or not p?"; answer options: "yes", "no", or "unclear"), and second, offered to those who initially provided a positive answer, where subjects could choose from two options: (i) "The protagonist knows that p"; (ii) "The protagonist feels like she/he knows that p, but she/he doesn't actually

know it". The researchers report that in those cases they investigated, only 41% of subjects deny knowledge in their first response, 24% deny it in the second, and only 35% make an "unwavering" knowledge attribution.[7] They also separately present data obtained for one Authentic Evidence case included their study, the *American Car* case, which was first used in experimental philosophy by Weinberg, Nichols, & Stich (2001),[8] later also examined in studies carried out by Cullen (2010). The scenario runs as follows:

> Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car (Nagel et al., 2013a, p. 660).

Only 38% of their participants denied knowledge on the first given occasion, with another 21% denying it when asked for the second time, while 41% attributed knowledge in both questions.[9] The results reported by Nagel et al. (2013a) show that a considerable number of laypersons who ascribe knowledge in Gettier-style cases revoke their judgment when given an opportunity and indicate that their antecedent attribution was a result of protagonist projection.

Machery et al. (2017a) also employed an analogous strategy as had Nagel et al. (2013a) in their experiment and observed a similar pattern of results—presenting a second, projective question regarding knowledge, resulted in a higher number of knowledge denials. For example, in the *Hospital* case mentioned in the previous section, 86% of Americans denied knowledge in the second question, while only 70% did so in the first question. The data for other nationalities show a similar effect—Brazil: 95% vs. 86%; Japan: 91% vs. 79%; India: 87% vs. 37%[10].

Interestingly, the level of protagonist projection for Indian subjects was considerably higher than for other nationalities, and offering them a projective question showed that their epistemic intuitions regarding the investigated Gettier-style cases do not differ in an important way from other ethnic groups.

The two studies discussed above suggest that the protagonist projection hypothesis might be a useful tool when it comes to explaining away unexpected folk knowledge attributions in Gettier-style cases. However, Machery et al. (2017a) study and the follow-up by Machery et al. (2017b) only focused on Apparent Evidence cases, while Nagel et al. (2013a) do not report precise data separately for Authentic and Apparent Evidence cases (with the small exception of the *American Car* scenario). I will now turn to present my experiment, in which I focus on knowledge attribution practices regarding five different Authentic Evidence cases.

4. The experiment

4.1. Research objectives

The main aim of the study was to estimate the level of protagonist projection in laypersons' reactions to different Authentic Evidence Gettier-style cases, or, to put it differently, measure how many knowledge attributions made by laypersons in such cases are only apparent, resulting from subjects' tendency to answer questions from the epistemic perspective of the protagonist in the presented scenario. How many subjects who initially ascribe knowledge in Authentic Evidence Gettier-style cases later withdraw the attribution when confronted with the paraphrasing, projective question?

4.2. Methods

The study materials included five vignettes: *American Car*, *Fake Barns*, *Thermometer*, *Wristwatch*, and *Bic Pen*. Each is based on a different Authentic Evidence case previously discussed in the literature (for the origins of each case—see sections 2 and 3 above). Among these vignettes, the *American Car* case was the only one that has been previously tested for protagonist projection (Nagel et al. 2013a) and was included in the experiment with intention of corroborating earlier findings.

The experiment was carried out as an online survey designed using an open-source software called LimeSurvey (www.limesurvey.org) and was posted on servers owned by KogniLab, X-Phi lab based at the University of Warsaw (www.kognilab.pl).

Each survey started with a demographic section which collected information on subjects' gender, age, and education. It also included two screening questions: one asked whether the participant was a native English speaker, the other explored whether the respondent studied philosophy in the past. Those subjects who admitted to not being native English speakers or reported having a degree in philosophy (BA, MA or PhD) were excluded from further statistical analysis.[11]

The experimental design was a between-subjects 2x5 factorial model, with the two following factors: Gettierization (control, test/Gettierized), and scenario (*American Car*, *Fake Barns*, *Thermometer*, *Wristwatch*, *Bic Pen*). Each participant was randomly assigned to one of the five vignettes and one of the two conditions. The conditions did not differ in the contents of the presented scenarios, but only in the questions asked to the participants: in the control condition the crucial knowledge question regarded a belief which was, in the considered situation, clearly an instance of knowledge, while in the test condition, the belief in question was

"Gettierized" (a true, justified belief that, according to intuitions shared by many epistemologists, does not count as knowledge).

After filling in the demographic section, each participant was presented with one vignette and answered a series of questions about it. First, they had to pass a short comprehension check consisting of two queries about some obvious factual features of the presented vignette. For the sake of illustration, in the *American Car* case they were as follows: C1) "Has Jill got an American car?"; answer options: "yes" or "no"; C2) "What kind of car has Jill got?"; answer options: "A Buick" or "A Pontiac". Respondents who failed to provide correct answers to these questions were not included in the final statistical analysis. Afterwards, the subjects were asked the crucial knowledge question—"Does [protagonist's name] know that [the target proposition]?"—and offered a 7-point answer scale ranging from "Doesn't know" (numbered 0) to "Knows" (numbered 6); only the endpoints were labeled. For the *American Car* case it was: K1-control) "Does Bob know that Jill has driven a Buick for many years?"; (K1-test) "Does Bob know that Jill drives an American Car?". Finally, the protagonist projection question was presented on a separate screen (so that the participants had no opportunity to change their antecedent judgments under the influence of the protagonist projection question).[12] The same format was employed for every vignette and condition: (K2) "What best describes the situation portrayed in the scenario?"; and the participants could choose from two options: (A) "[protagonist's name] thinks (s)he knows that [the target proposition], but (s)he doesn't really know that"; (B) "[protagonist's name] really knows that [the target proposition]".[13]

4.3. Participants

A total of 891 respondents participated in the study. All were recruited via a professional respondent panel, ClickWorker (www.clickworker.com), and received a small financial compensation for filling in the survey. The exclusion procedure (see section 4.2. for details) yielded a final sample size of 805. Of those 805 subjects, 166 were assigned to the *American Ca*r vignette, 171 assessed the *Fake Barns* case, 161 the *Thermometer* case, 153 the *Wristwatch* case, and 154 evaluated the *Bic Pen* scenario.[14] 62% of the participants were female. The age distribution ranged from 18 to 93 ($M = 35.02$; $SD = 11.47$), and was slightly skewed (50% of the respondents were 33 years old or younger).[15] When it comes to education, the sample was quite diverse: 36% of subjects had a Bachelor's degree, 11% a Master's degree, 2% a PhD, 32% underwent some college education but did not receive a degree (or have not received one yet), and 16% did not attend college at all (the remaining 3% indicated the option 'other').

4.4. Results

In the presentation of the data I will first focus on participants' answers to the initial knowledge question (K1) depending on Gettierization (control vs. test condition) and scenario. In order to analyze the data concerning K1, I employed a 2x5 between-subjects ANOVA analysis.

A significant main effect of Gettierization was observed—the participants were, on average, less likely to attribute knowledge in the test condition ($M = 4.83$, $SD = 2.04$) than in the control condition ($M = 5.71$, $SD = 0.95$); $F(1,795) = 64.79$; $p < 0.001$; $\eta^2 = 0.075$. The analysis revealed a significant main effect of scenario as well: $F(4,795) = 15.98$; $p < 0.001$; $\eta^2 = 0.074$. Further *post-hoc* comparisons between scenarios (Bonferroni method, $p < 0.05$) show that this followed from the fact that, on average, the respondents were less happy to attribute knowledge

in the *Fake Barns* case (M = 5.05, SD = 1.82) than in the *Wristwatch*, and *Bic Pen* cases. Also, their tendency to ascribe knowledge was significantly lower in the *American Car* case (M = 4.58, SD = 2.13) compared to all other four scenarios. There were no differences between the *Thermometer* case (M = 5.43, SD = 1.46), the *Wristwatch* case (M = 5.71, SD = 1.06), and the *Bic Pen* case (M = 5.66, SD = 1.15). Moreover, a significant interaction effect between the two factors emerged: $F(4,795)$ = 5.6; p < 0.001; η^2 = 0.027. Pairwise comparisons between the control and test condition performed independently for each scenario (Bonferroni method, p < 0.05) revealed that this interaction appeared because Gettierization had the predicted impact on knowledge attributions only in the *American Car*, *Fake Barns*, and *Thermometer* cases, but not in the *Wristwatch* and *Bic Pen* cases. The average knowledge ratings for each scenario depending on Gettierization are illustrated in Figure 1.

The data corroborates previous experimental findings. As in studies by Colaço et al. (2014) and Ziółkowski (2016), I found a significant effect of Gettierization in the *Fake Barns* case (d = 0.75, which could be classified as mid-to-large effect). Moreover, a similar effect, but smaller in size (d = 0.57), was observed for the *Thermometer* case, which supports data presented by Ziółkowski (2016). The size of the Gettierization effect was the most prominent in the case of the *American Car* scenario (d = 0.79), which is in line with data collected by Cullen (2010) and Nagel et al. (2013a)—although they used a different, "yes-or-no" answer scale and did not include a control condition in their experiments, they observed that many laypersons tend to deny knowledge in the Gettierized condition when confronted with the *American Car* scenario. And lastly, high knowledge ratings for the *Wristwatch* and *Bic Pen* cases fit the same pattern as the findings reported by Starmans & Friedman (2012). It is worth noting that all three Gettierization effects I observed were not very large and the average knowledge ratings were

clearly above the midpoint of the scale in both the control and test condition for each investigated scenario, which means that laypersons were rather attributing knowledge than denying it across all experimental conditions (for measure K1).
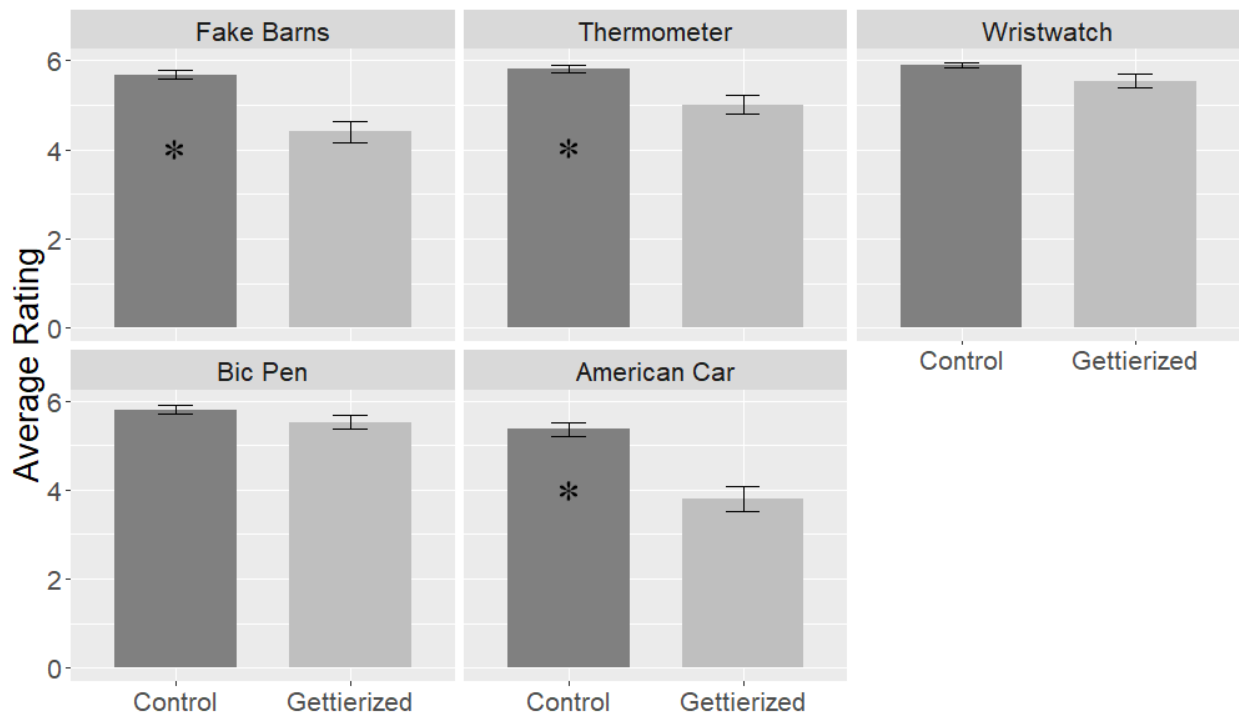

Figure 1. Average answers to the first knowledge question (K1) in each experimental condition. 0 = 'doesn't know'; 6 = 'knows'. An asterisk indicates a significant difference between ratings in the control and test condition (Bonferroni *post-hoc*, $p < 0.05$), no asterisk means lack of statistical significance.

I now turn to present the data collected in the projective knowledge question (K2), where the participants had a chance to withdraw their initial knowledge attributions and admit projecting the perspective of the protagonist. Figure 2. below presents the percentage of

participants who attributed knowledge in K1 (judgments above the midpoint of the scale) and K2 in each scenario and experimental condition.
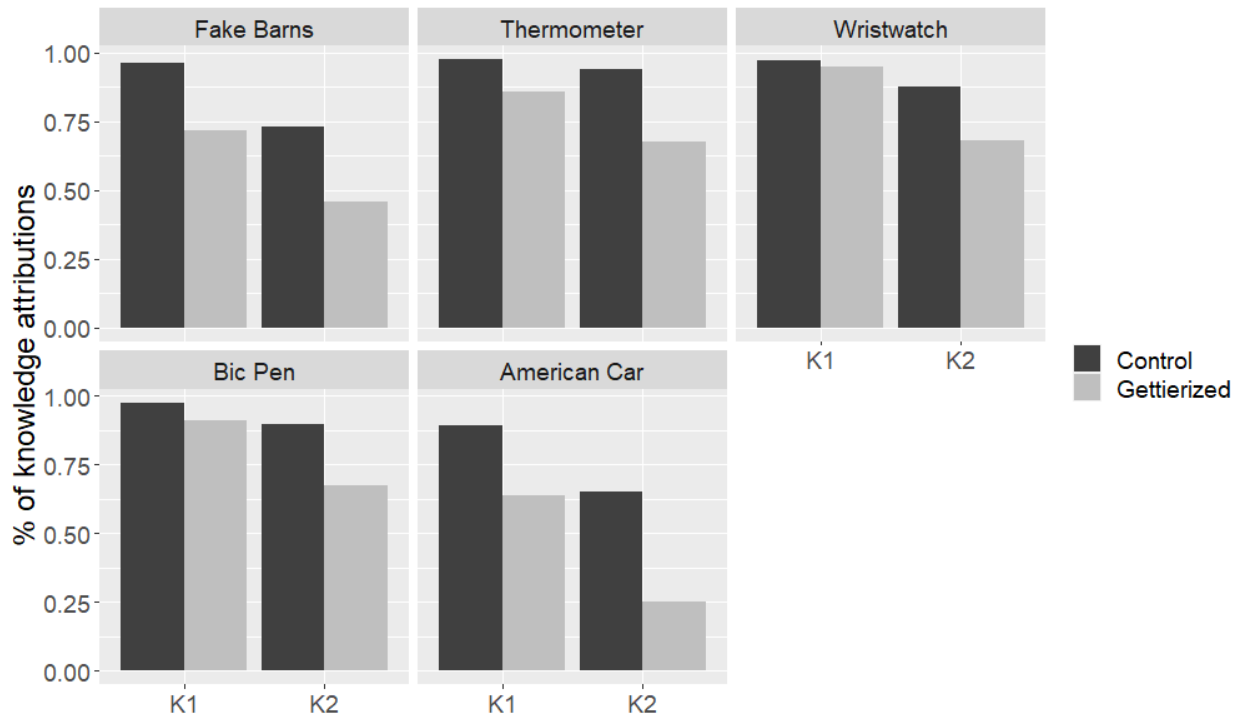

Figure 2. The proportion of laypersons attributing knowledge in K1 (Likert scale rating above the midpoint) and K2 (two-point projective question).

Across all five examined Authentic Evidence cases, the majority of respondents initially ascribed knowledge in the Gettierized condition, i.e., provided a positive answer to the first knowledge question (K1), although for the *American Car* and *Fake Barns* scenarios the number of attributions was considerably lower. Importantly, in the *Thermometer*, *Wristwatch*, and *Bic Pen* cases, the majority of participants still attributed knowledge, even when asked about it for the second time in the projective question (K2). In the *Fake Barns* case subjects' answers to K2 split almost in half. The *American Car* case stands out from the scenarios tested in the

experiment, as it was the only vignette that encouraged most of the participants (75%) to deny knowledge in K2.

In order to investigate which factors affected subjects' judgments in K2, I performed a logistic regression analysis[16]. The first model included Gettierization and scenario as factors. The model was statistically significant (χ^2(5) = 150.2; $p < 0.001$) and explained 23.8% of variance (Nagelkerke R^2) with 74% level of accuracy. Unsurprisingly, a significant and considerably strong main effect of Gettierization emerged—subjects were less likely to attribute knowledge in K2 in the Gettierized condition than in the control condition ($p < 0.001$). Scenario also played a significant role here—the baseline for comparisons was the *American Car* condition, where the number of knowledge denials was the most prominent. It turned out that, comparing to this case, participants were considerably more likely to attribute knowledge in K2 for the *Bic Pen*, *Wristwatch*, and *Thermometer* scenarios ($p < 0.001$; with similar *B* coefficients) and slightly more likely to do it in the *Fake Barns* case ($p < 0.01$). Detailed results of the analysis can be found in Tables 1. and 2. in Appendix B. I also tested another Logit model which included interactions between Gettierization and scenario to see whether the impact of Gettierization on K2 judgments was different between scenarios. None of the interactions turned out to be significant, and overall the model without interactions fit the data better than the model including interactions (see Appendix B). Thus, although some scenarios elicited more knowledge denials in K2 than others, we have no reason to say that Gettierization affected K2 judgments differently across scenarios.
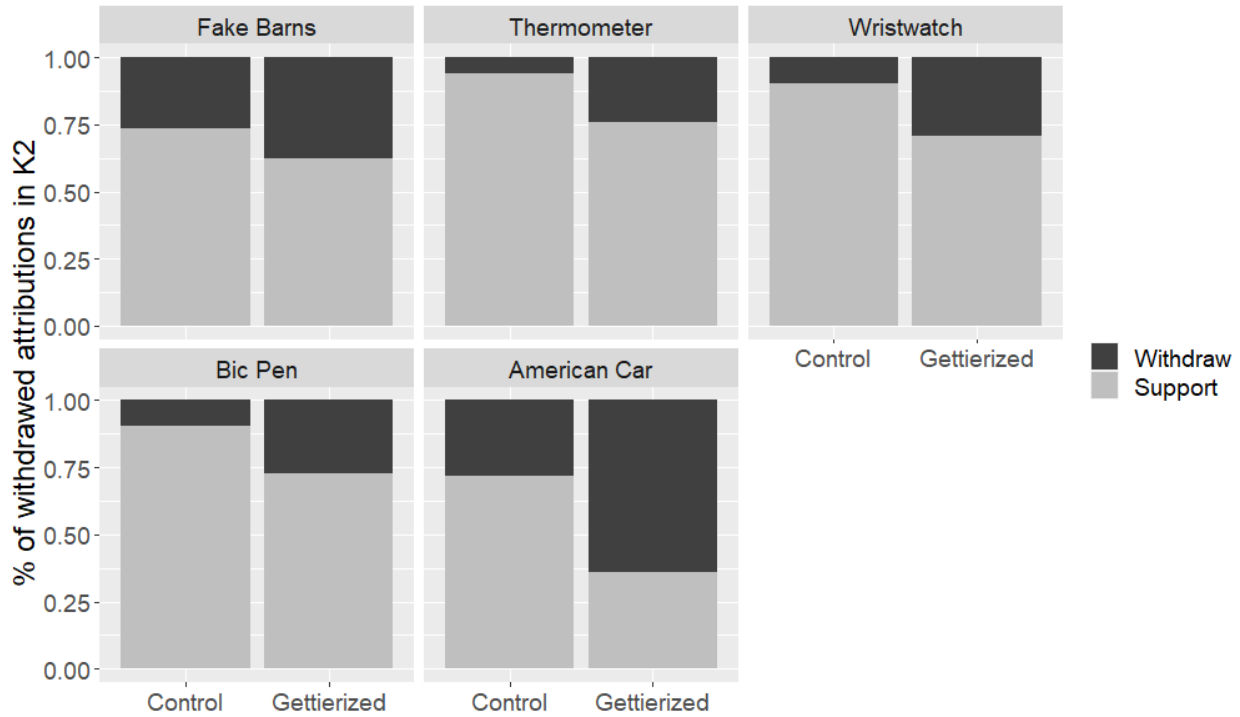
Figure 3. The percentage of subjects who first attributed knowledge in K1 and later withdrew or supported their attribution in K2.

Figure 3. presents the relation between positive ratings (knowledge attributions) in K1 and final judgments in K2. Across all the scenarios and conditions, 711 of 805 subjects (88.3%) attributed knowledge in K1; the figure shows how many of those subjects later withdrew or supported their attribution when confronted with the projective question (K2). If we observed that the majority (or at least a noticeable portion) of participants revoke their initial attributions when offered a projective question, then we could conclude the data strongly support the protagonist projection hypothesis. But that's not what we obtained for the *Thermometer*, *Wristwatch*, and *Bic Pen* cases, where the majority of participants supported their initial attribution[17]—in the Gettierized condition less than one third of the respondents who first attributed knowledge admitted projecting, which suggests that protagonist projection was not the

main reason why laypersons ascribe knowledge in these cases. Compared to the three abovementioned cases, the number of subjects withdrawing their initial attribution in the test condition for the *Fake Barns* case was somewhat higher (38%). Interestingly, here a substantial number of subjects (27%) also changed their mind in the control condition, which was supposed to be a clear case of knowledge. Once again, the *American Car* scenario stood out from the rest, since here most of the respondents (64%) who first ascribed knowledge in the test condition later denied it[18]. However, similarly as in the case of *Fake Barns* scenario, many subjects (28%) decided to resign from their antecedent knowledge attribution in the control condition, which was not expected.

5. Discussion and conclusions

The results obtained in the experiment discussed above are twofold—they confirm our expectations based on previous empirical findings as far as folk knowledge attribution practices in Authentic Evidence Gettier-style cases are concerned, but, at the same time, they do not fully support our hopes regarding the role of protagonist projection in the mechanics of these attributions.

I will present two alternative interpretations of the data. What needs to be noted, though, is that these interpretations are mutually exclusive, which I am aware of. The first evaluates the protagonist projection predictions based on the data and assumes that our measure of protagonist projections is accurate. The second interpretation is slightly more speculative and addresses some doubts regarding this latter assumption—maybe the traditional methods do not allow us to precisely diagnose the level of protagonist projection.

5. 1. Taking data at face value

If the protagonist projection hypothesis played the expected role in explaining away problematic folk knowledge attributions in Authentic Evidence cases, we should have observed that a large portion (preferably: majority) of subjects who initially attribute knowledge (K1) later withdraw their attributions when presented with the projective question (K2). But that did not happen for the *Thermometer*, *Wristwatch*, and *Bic Pen* cases. The data does not support the prediction that the inclination of laypersons to attribute knowledge in these scenarios is due to protagonist projection. On the other hand, the protagonist projection hypothesis did bear out for the *American Car* and—to some degree—*Fake Barns* cases.

If we aggregate the data from all five scenarios, it turns out that in the Gettierized condition, out of those subjects who initially ascribed knowledge ($n = 325$) only 35% later denied it when offered an opportunity to choose the projective answer. It means that the majority of laypersons (65%) who attribute knowledge in Authentic Evidence cases included in the experiment uphold their attributions and deny that they intended to project the perspective of the protagonist onto their answers. The *American Car* case remains the only one in which the majority of participants decided to resign from their initial positive knowledge judgments. Here, 75% of laypersons denied knowledge when presented with the projective question (K2) compared to 64% of initial knowledge ascriptions (K1). In *Fake Barns* answers to the projective question (K2) in the Gettierized condition split almost in half (compared to 72% of initial attributions), which might suggest some projection at play, but still leaves us with a substantial number of unexplained folk knowledge ascriptions.

The results might seem odd at first sight, but we should take into account that, at least when initial knowledge attributions are concerned, they are in line with previous findings I briefly discussed in Sections 2 and 3. However, noticeable differences between scenarios shown by the data require some explanation. Interestingly, the pattern here is similar for both knowledge measures (K1 and K2)—on average subjects were less likely to attribute knowledge in the *American Car* and *Fake Barns* cases than in the other scenarios investigated in the experiment (this difference was slightly more pronounced for K2 than K1).

Blouw, Buckwalter, & Turri (2017) offered an interesting explanatory framework that could be helpful here—they present a taxonomy of Gettier-style cases which tracks differences between scenarios along three dimensions: *Detection*, *Threat*, and *Replacement*.[19] *Detection* concerns whether or not an agent initially succeeds or fails in detecting the truth value of the believed proposition and it is useful in drawing a distinction between Apparent Evidence and Authentic Evidence cases—while agents fail to detect certain falsehoods in the former type, there is no such failure in the latter type. *Threat* tracks the presence of epistemic misfortune and luck—one can form a true belief that might be disrupted by misfortune, but is not (threat is *unsuccessful*) or a true belief whose truth-making relation ends up being disrupted by a stroke of bad luck (threat is *successful*). *Replacement* concerns those cases where *Threat* is successful, but thanks to another stroke of luck the truthmaker of the belief in question is being replaced in the process—the replacement might be very *similar* to the original truthmaker, or quite *dissimilar*. While all five vignettes investigated in my experiment seem to be alike as far as *Detection* is concerned, there are some significant differences along the *Threat* and *Replacement* dimensions. While *Threat* is unsuccessful in the *Fake Barns* and *Thermometer* cases, it is successful in the other three. As a result, *Replacement* plays no important role for *Fake Barns* and *Thermometer*

(the truth-making relation is not broken), but it does for the other three. In the *Wristwatch* and *Bic Pen* cases the replaced truthmaker is very similar to the original (qualitatively identical watch or pen), but in the *American Car* scenario it is quite dissimilar (a Pontiac is a different kind of American car than a Buick). According to the model offered by Blouw et al., we should expect the following ordering of cases based on the proportion of knowledge attributions: *Failed Threat > Successful Threat + Similar Replacement > Successful Threat + Dissimilar Replacement*.

This framework might explain why the *American Car* scenario (*Dissimilar Replacement*) elicited the smallest number of knowledge ascriptions, and, in particular, smaller from ratings for the *Wristwatch* and *Bic Pen* cases (*Similar Replacement*). Relatively high knowledge ratings for the *Thermometer* case (*Failed Threat*) also fit this model. However, a substantial number of subjects denied knowledge in the *Fake Barns* scenario (also a *Failed Threat* case), which is the opposite of what the model predicts. The fact that the *Thermometer* case seems to be on a par with the *Wristwatch* and *Bic Pen* cases is another thing that lowers the explanatory values of the framework offered by Blouw, Buckwalter, & Turri (2017). Thus, some differences between the intuitions evoked by the scenarios investigated in my experiment remain unclear. Since exploring those differences was not the main aim of this study, I will just express my hope that future research will shed more light on these aspects of my findings.

5. 2. Methodological reservations: do we have a good measure of protagonist projection?

Unfortunately, besides some results seemingly favorable to the protagonist projection hypothesis, the projective knowledge question (K2) yielded some puzzling data as well. What is extremely surprising, in those cases where a significant number of subjects found the projective answer attractive in the Gettierized condition (the *American Car* and *Fake Barns* cases), many subjects

also tended to withdraw their preceding knowledge attribution in the control condition (28% and 27% respectively), where no such tendency was predicted. These ratios are comparable to the ratios of knowledge attribution withdrawals for the *Thermometer*, *Wristwatch*, and *Bic Pen* cases, but in the Gettierized condition. It is difficult to find a good reason for denying knowledge in the control condition (regardless of which scenario we consider), since in this situation the respondents were asked about beliefs which are non-controversial instances of knowledge.[20] One might hypothesize that laypersons are simply being "pushed around" here—if the researcher insistently keeps asking about knowledge and offers another opportunity to deny it, the respondents take it to be an explicit suggestion which they choose to follow. That would mean that not all knowledge withdrawals in K2 in fact indicate protagonist projection, but maybe some other pragmatic mechanism (one might refer to the Gricean framework here). However, this explanation cannot account for the fact that almost no participants withdrew their antecedent knowledge ascriptions in the control condition when they were confronted with the *Thermometer*, *Wristwatch*, and *Bic Pen* scenarios. Thus, the result in question remains baffling and might actually hint that the paraphrasing question (K2) I adopted from the studies by Nagel et al. (2013a) and Machery et al. (2017a, 2017b) is not a fully accurate measure of protagonist projection.

Similar worries were recently addressed by Gonnerman et al. (2020) who investigated whether protagonist projection might be responsible for the inclination of laypersons to classify some unjustified beliefs as knowledge. While their data show some support for the protagonist projection hypothesis in this respect, at the same time they suggest that not all seemingly projective answers provided by subjects stem from protagonist projection—at least some of them might result from some other cognitive mechanisms.

Although the paraphrasing question constructed along the lines of my K2 prompt has a long tradition of usage in X-Phi research, its validity has never been evaluated. In the light of results such as those obtained by Gonnerman et. al (2020) or in my study, developing an alternative measure of protagonist projection (even for the purpose of cross-validation) might be a worthwhile research enterprise.

Appendix A: survey questions

Fake Barns—survey questions:

C1 (control/test): "Does Gerald think he saw a [cow / house]?" (answer options: "yes" or "no").

C2 (control/test): "Did Gerald see a [cow / house]?" (answer options: "yes" or "no").

K1 (control/test): "Does Gerald know he saw a [cow / house]?"

K2 answer options (control/test): "Gerald thinks he knows that he saw a [cow / house], but he doesn't know that" or "Gerald really knows that he saw a [cow / house]".

Thermometer—survey questions:

C1 (control/test): "Does Stephen think his [son has red eyes / son's temperature is 98.6]?" (answer options: "yes" or "no").

C2 (control/test): "[Does his son have red eyes / Is his son's temperature 98.6]?" (answer options: "yes" or "no").

K1 (control/test): "Does Stephen know his [son has red eyes / son's temperature is 98.6]?"

K2 answer options (control/test): "Stephen thinks he knows that his [son has red eyes / son's temperature is 98.6], but he doesn't know that" or "Stephen really knows that his [son has red eyes / son's temperature is 98.6]".

Wristwatch—survey questions

C1 (control/test): "Is there a [book / watch] on the table?" (answer options: "yes" or "no").

C2 (control/test): "How did the [book / watch] get on the table?" (answer options: "The burglar put it there" or "Peter put it there").

K1 (control/test): "Does Peter know that there is a [book / watch] on the table?"

K2 answer options (control/test): "Peter thinks he knows that there is a [book / watch] on the table, but he doesn't know that" or "Peter really knows that there is a [book / watch] on the table".

Bic Pen—the vignette and survey questions:

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 on the coffee table. Then the burglars leave. Katie is still in the shower, and did not hear anything (Starmans & Friedman, 2012, p. 276).

C1 (control/test): "Is there a [letter / blue Bic pen] on the table?" (answer options: "yes" or "no").

C2 (control/test): "How did the [letter / blue Bic pen] get on the table?" (answer options: "The burglar put it there" or "Katie put it there").

K1 (control/test): "Does Katie know that there is a [letter / blue Bic pen] on the table?"

K2 answer options (control/test): "Katie thinks she knows that there is a [letter / blue Bic pen] on the table, but she doesn't know that" or "Katie really knows that there is a [letter / blue Bic pen] on the table".

Appendix B: **results of logistic regression for K2 judgments**

Table 1. Logit model for K2 responses without interactions

Predictor	*B*	*SE B*	*OR*	*OR 95% CI*	*p*-value
Condition					
Control	—	—	—	—	—
Gettierized	-1.47	0.18	0.23	[0.16; 0.32]	<0.001
Scenario					
American Car	—	—	—	—	—
Bic Pen	1.67	0.27	5.34	[3.19; 9.11]	<0.001
Fake Barns	0.66	0.24	1.94	[1.22; 3.09]	0.005
Thermometer	1.83	0.27	6.24	[3.71; 10.8]	<0.001
Wristwatch	1.65	0.27	5.20	[3.11; 8.84]	<0.001
Constant	0.52	0.19			0.005

AIC: 870.19

Table 2. Logit model for K2 responses including interactions

Predictor	*B*	*SE B*	*OR*	*OR* 95% CI	p-value
Condition					
Control	—	—	—	—	—
Gettierized	-1.70	0.34	0.18	[0.09; 0.35]	<0.001
Scenario					
American Car	—	—	—	—	—
Bic Pen	1.53	0.44	4.63	[2.04; 11.6]	<0.001
Fake Barns	0.39	0.34	1.47	[0.76; 2.86]	0.25
Thermometer	2.14	0.52	8.49	[3.34; 26.2]	<0.001
Wristwatch	1.36	0.42	3.88	[1.75; 9.34]	0.001
Condition × Scenario					
American Car – Gettierization	—	—	—	—	—
Bic Pen – Gettierization	0.28	0.56	1.33	[0.43; 3.92]	0.62
Fake Barns – Gettierization	0.53	0.47	1.70	[0.67; 4.32]	0.26
Thermometer – Gettierization	-0.32	0.62	0.72	[0.20; 2.35]	0.60
Wristwatch – Gettierization	0.50	0.55	1.64	[0.55; 4.77]	0.37
Constant	0.62	0.23			0.007

AIC: 875.27

NOTES

[2] Nagel et al.(2013b) argue that the distinction between Apparent and Authentic Evidence cases is far from being well-defined and clear-cut and that there are several possible understandings of it. For the sake of this paper, I propose to employ the understanding fleshed out at the beginning of this section. In Apparent Evidence cases the belief which aspires to be knowledge follows from another belief which was false at the moment of its formation (e.g., Smith believes that Jones owns a Ford or Brown is in Barcelona because Smith has a false belief that Jones owns a Ford).

[3] For a detailed discussion concerning the role of luck for knowledge attributions and distinctions between different kinds of epistemic luck, see for example Unger (1968) and Pritchard (2005).

[4] This small difference did not pass the threshold of statistical significance, although the researchers observed a significant effect of Gettierization on the composite score, where they

combined knowledge attributions with partcipants' confidence ratings. I am grateful to an anonymous Referee for pointing this out.

[5] The study also included other purported counterexamples to the factivity assumption containing other verbs taken to be factive.

[6] The results were considerably less favorable to the protagonist projection hypothesis for 'learn' and 'realize'—in these cases about half of subjects gave non-projective answers (which are inconsistent with the factivity assumption).

[7] Nagel et al. (2013a) included few different Apparent and Authentic Evidence Gettier-style cases in their study, but they report aggregated results computed together for all the vignettes. They also remark that subjects were 2.28 times more likely to immediately ascribe knowledge in the Authentic Evidence cases than in the Apparent Evidence cases (reactions to the first knowledge question), but do not report any data on possible differences between reactions to the second knowledge question depending on the type of case.

[8] Weinberg, Nichols, & Stich (2001) reported data that famously suggested the existence of cross-cultural differences in folk intuitions regarding the *American Car* case and some other important epistemological thought experiments. However, recent attempts to replicate their controversial discoveries were unsuccessful (see Seyedsayamdost, 2015; Kim & Yuan, 2015). I will not discuss Weinberg, Nichols, & Stich's (2001) findings here, since their data is dubious.

[9] In his earlier study, Cullen (2010) also presented his respondents with the *American Car* case and asked them only one question: "Does Bob know that Jill drives an American Car?"; answer options: "knows", "does not know". He found that 58% of his participants denied knowledge right away (compared to 38% in Nagel et al. 2013a study). One might suspect that if the

projective question was included in his experiment, the observed number of problematic

knowledge ascriptions would be even smaller than found by Nagel et al. (2013a).

[10] It is important to note that these statistics concern only answers provided by participants who

judged that the belief in question was justified. Machery et al. (2017a) employed this procedure

to exclude subjects who did not view the case as a genuine Gettier-style case (i.e. found the

belief in question not justified). I am grateful to an anonymous Referee for pointing out that

issue.

[11] I believe that subjects who received a degree in philosophy most likely had contact with some

(if not all) scenarios used in the experiment. Since we are interested in untutored and untrained

intuitions, I think that their answers should not be taken into consideration.

[12] The vignette itself remained at the top of the screen, giving subjects the opportunity to consult

the contents of the scenario when they were answering the second knowledge question.

[13] The exact formulation of questions for the *Fake Barns*, *Thermometer*, *Wristwatch*, and *Bic Pen*

cases can be found in Appendix A.

[14] These small differences in the size of groups were a result of the employed screening

procedure, which excluded slightly different numbers of participants depending on the vignette

and experimental condition.

[15] Two respondents did not provide valid information about their age.

[16] I am grateful to the Editors of Oxford Studies in Experimental Philosophy, vol. 4, for

suggesting this approach.

[17] Chi-square tests aimed at establishing whether the proportion of supported knowledge

attributions exceeded 0.5—*Thermometer*: $\chi^2(1) = 15.5$; $p < 0.001$; *Wristwatch*: $\chi^2(1) = 12.8$; $p <$

0.001; *Bic Pen*: $\chi^2(1) = 14.6$; $p < 0.001$.

[18] $\chi^2(1) = 4.25$; $p = 0.039$.

[19] I am grateful to an anonymous Referee for helpful comments which lead to substantial improvements in this passage.

[20] Since the *Fake Barns* case strongly involves perception in the process of belief acquisition (in both the control and test condition), and perception is widely believed to be a fallible cognitive mechanism, we might expect laypersons to be uncertain about their knowledge attributions regarding this scenario. After all, Gerald looks at objects he describes (cows, houses) from a distance and on the move—clearly, sight might be faulty in such circumstances. Asking the second knowledge question (K2) might have made this possibility of error more salient to the participants and lead them to withdraw their antecedent knowledge ascriptions even in the control condition. Unfortunately an analogous line of reasoning cannot be used to explained away unpredicted knowledge attribution withdrawals observed for the *American Car* case, where perception does not play a similar role in belief acquisition. I am grateful to Mieszko Tałasiewicz for this suggestion.

REFERENCES

Blouw, P., Buckwalter, W., & Turri, J. (2017). Gettier Cases: A Taxonomy. In R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining Knowledge: New Essays on the Gettier Problem,* Oxford: Oxford University Press.

Buckwalter W. (2014). Factive Verbs and Protagonist Projection. *Episteme*, 11 (4), 391–409.

Cohen, J. (1988). *Statistical power analysis for behavioral sciences* (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum.

Colaço, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic Intuitions in Fake-Barn Thought Experiments. *Episteme*, 11 (2), 199–212.

Cullen, S. (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1 (2), 275–296.

Gettier, E. (1963). Is Justified True Belief Knowledge?. *Analysis*, 23 (6), 121–123.

Gonnerman, C., Poag, L., Redden, L., Robbins, J. & Crowley, S. (2020). In Our Shoes or the Protagonist's? Knowledge, Justification, and Projection. In: T. Lombrozo, S. Nichols & J. Knobe (eds.), Oxford Studies in Experimental Philosophy, Vol. 3., Oxford: Oxford University Press.

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

Goldman, A. I. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Holton, R. (1997). Some Telling Examples: A Reply to Tsohatzidis. *Journal of Pragmatics*, 28, 624–628.

Kim M., & Yuan Y. (2015). No Cross-Cultural Differences in the Gettier Car Case Intuition: A Replication study of Weinberg et al. 2001. *Episteme*, 12 (3), 355–361.

Machery E., Stich, S., Rose D., Chatterjee A., Karasawa K., Struchiner N., Sirker S., Usui N., & Hashimoto T. (2017a). Gettier Across Cultures. *Noûs*, 51 (3), 645–664.

Machery, E. et al. (2017b). The Gettier Intuition from South America to Asia. *Journal of Indian Council of Philosophical Research*, 34 (3), 517–541.

Nagel, J., San Juan, V., & Mar, R. (2013a). Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, 129 (3), 652–661.

Nagel, J., San Juan, V., & Mar, R. (2013b). Authentic Gettier Cases: a reply to Starmans and Friedman. *Cognition*, 129 (3), 666–669.

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

Seyedsayamdost H. (2015). On Normativity and Epistemic Intuitions: Failure of Replication. *Episteme*, 12 (1), 95–116.

Starmans, C., Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124 (3), 272–283.

Unger, P. (1968). An Analysis of Factual Knowledge. *Journal of Philosophy*, 65, 157–170.

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

Ziółkowski, A. (2016). Folk Intuitions and the No-Luck-Thesis. *Episteme*, 13 (3), 343–358.