

Gettier and Context

Delia Belleri (University of Vienna), Annalisa Coliva (University of California, Irvine)

Abstract

The relationship between contextualism and Gettier cases is controversial. Yet, David Lewis, in his influential “Elusive Knowledge”, links his contextualist thoughts to the discussion of some standard Gettier cases. This chapter explores the question whether contextualism can provide a satisfactory account of Gettier scenarios. We compare a contextualist account of lottery cases and of fake-barn cases (which Lewis seems to regard as very close to Gettier examples) with a contextualist account of standard Gettier problems. We conclude that contextualism may at best provide a partial reconstruction and explanation of the problem underlying Gettier scenarios.

1. Introduction

Epistemic Contextualism is known as the view that the truth-value of knowledge ascriptions depends on the epistemic standards operative in the attributor’s context (Cohen 1998, 1999; DeRose 1992, 1995, 1999; Heller 1999; Lewis, 1996). The higher the standards, the easier it is for a knowledge ascription “S knows that P” to be false. Typically, when standards are high, the subject’s evidence is unable to rule out some *relevant alternatives* where the proposition p under discussion is false (Dretske 1981; Goldman 1976).

Contextualism has been proposed primarily as a way of dealing with skeptical problems, such as evil demon-, or brain in a vat-scenarios (Cohen 1999, DeRose 1995, Neta 2003). It also purports to explain the apparent contextual variability of “know”-statements in everyday contexts, where practical stakes seem to affect our epistemic position: Cohen’s (1999) “airport case” and DeRose’s (1992) “bank case” provide apt illustrations of the close relationship between practical stakes and the truth-value of knowledge ascriptions.

Contextualism seems nicely to explain how and why contextual factors, such as high epistemic standards, can “destroy” knowledge. Skeptical and certain high-stakes practical worries are certainly good examples of this phenomenon, but they are not the only ones. Three further examples in which knowledge is somehow “destroyed” will be the focus of the

present paper: “fake-barn” cases, “lottery” cases and Gettier cases. Let us briefly review them one by one.

(1) *Fake Barn Cases*: Suppose Henry is driving through a country that is replete with barn façades, and only contains a few real barns. Henry is not privy to this situation. He looks at what appears to him as a barn and comes to believe that he is seeing a barn. Indeed, the barn he is looking at is a real one, but given how few real barns there are in the country, he could have easily been looking at a fake one. In such a scenario, it would seem false to say that Henry knows that he is seeing a real barn (Goldman 1976).

(2) *Lottery Cases*: John has just bought a ticket of a fair lottery. There is a 99,99% chance that his ticket will lose, plus, suppose his ticket is indeed not the lucky one. John declares: “I know I will loose the lottery”. Despite there being an extremely high chance for this to happen, John’s self ascription of knowledge seems nevertheless wrong, because for all John knows, his ticket could be the lucky one (see Hawthorne 2005).

(3) *Gettier Cases*: Smith’s wrist-clock has stopped, even though Smith hasn’t realized it yet. Smith looks at his clock, which tells him it is 8:29, so Smith comes to believe it is 8:29. As it happens, it is 8:29, so Smith’s belief is true. Yet, it does not amount to knowledge. (Gettier, 1963 is obviously credited with the first Gettier cases ever appeared; the stopped clock case is owed to Russell, 1948).

Does contextualism offer a good account of cases like (1), (2) and (3), too? In his influential 1996 paper “Elusive Knowledge”, David Lewis answers in the positive to this question. He links fake-barn cases and lottery cases to Gettier cases, *plus* he argues that his brand of contextualism can account for Gettier cases.

In this paper, we wish to argue that Lewis’s view has only a limited reach. Not only it is contentious to argue that fake-barn and lottery cases are closely linked to Gettier cases as

Lewis suggests; it is generally objectionable to regard contextualism itself as a good framework for thinking about Gettier cases. In what follows, we first reconstruct Lewis' contextualist take on Gettier cases, fake barns and lottery cases. Secondly, we compare Gettier cases with lottery cases. Thirdly, we compare Gettier cases with fake-barn cases. Lastly, we point at the limitations of contextualism as an account of Gettier cases.

2. Lewis' Contextualism about Epistemic Predicates

Lewis introduces Contextualism as a way of escaping the skepticism that threatens an infallibilist account of knowledge. In order to avoid the pitfalls of skepticism, Lewis opts for a qualified form of infallibilism, whereby a subject knows that p if and only if p obtains in a specific set of possibilities, namely all the possibilities left uneliminated by her evidence. Equivalently stated, S knows that p if and only if S's evidence *eliminates all non- p possibilities*.

The set of *uneliminated possibilities* is a set of possible worlds where the subject has the same experience and memories as those she actually has. Knowledge is infallible only relative to this restricted set of possibilities, insofar as in every scenario where I have that particular experience and memory, the proposition p I believe is true. Yet, what if among the uneliminated possibilities there were still scenarios where p is false? Lewis manages to deal with these threatening possibilities by allowing that they be *ignored* in the context of the knowledge attributor, for instance on the grounds that they are not salient. This delivers Lewis' final formulation of his brand of contextualism, which is as follows:

S knows that P iff S's evidence eliminates every possibility in which not-P - Psst! - except for those possibilities that we [attributor and hearer] are properly ignoring. (1996, p. 54).

Which criteria determine whether a possibility is or is not properly ignored? Lewis gives a list of rules to serve this purpose. We only need to mention two of them here: First, according to the *Rule of Actuality*, the possibility that happens to be actual is never properly ignored. Next, the *Rule of Resemblance* has it that if one possibility resembles another, and one of the two is not properly ignored, then neither is the other.

Let us now go back to the (1)-(3) cases listed earlier. Following Lewis, we may say that knowledge is “destroyed” because, from the perspective of the attributor, the subject violates one of the rules listed above, thus *improperly ignoring* certain possibilities. The following subsection reconstructs in closer detail Lewis’ account of (1)-(3).

2.2 Lewis on Gettier, Fake Barns and Lotteries: Improperly Ignored Possibilities

Lewis purports to diagnose and resolve Gettier cases by invoking his *Rule of Resemblance*, according to which if one possibility saliently resembles another, and one of the two is not properly ignored, then neither is the other.

Suppose Sue believes that either Nogot or Havit owns a Ford. She believes that Nogot owns a Ford on the basis of her having seen him driving a Ford, while she has no evidence suggesting that Havit is the owner of a Ford. Suppose, however, that Nogot owns no car at all: the vehicle he was driving when he was spotted by Sue was rented. Plus, Havit does own a Ford, although he almost never drives it and Sue never saw him doing so. Then Sue’s belief *that either Nogot or Havit owns a Ford* is true (because Havit owns a Ford) and justified (on the basis of Sue’s spotting Nogot on a Ford), yet it is not knowledge. Why does Sue fail to achieve knowledge in this case? According to Lewis, because – from the point of view of a knowledge-attributor – she *improperly ignores* a possibility that is still compatible with her evidence: namely, that Nogot drives a Ford he does not own, while Havit neither drives nor owns a car. Why is this possibility improperly ignored? Because, first, actuality is never properly ignored; and second, because this possibility *saliently resembles actuality* as far as Sue’s evidence is concerned, and since actuality is not properly ignored, neither is this resemblant possibility.

Notably, Lewis believes the same explanation holds for fake-barn cases and lottery cases. These respond to the same resemblance considerations that, in Lewis’ view, seem *central* to explain gettierization. To the extent that they contain these central ingredients too, it seems safe to say that he considers them as very close to Gettier cases.

Imagine Ida is traveling in a country full of fake barns and finds herself in the fake-barn situation described in (1). On the Lewisian account, Ida lacks knowledge in that she is improperly ignoring a possibility that her evidence leaves uneliminated: namely that she is looking at one of the many barn façades in the country. This possibility is not properly

ignored because it resembles actuality – relative to the evidence at her disposal –, too much to be set aside, given the amount of bogus barns in Ida’s surroundings.

Further support to the idea that fake-barn cases are closely linked to Gettier cases in Lewis’ sense seems to come from contrastivism (Schaffer 2004, 2005). Contrastivism has it that knowledge attributions generally have the form “S knows that *p* rather than *q*”. As Schaffer notes (2005, p. 243), contrastivism fits particularly well the way in which perception works, namely through *discrimination*. Sensory discrimination seems to enable us to reject some possibilities based on the inputs we receive. For instance, outside of the fake barns country, if Ida had the visual experience as of the façade of a building, this input, together with the reasonable assumption that she is in the real world and that façades are usually parts of whole buildings, would enable her to rule out the possibility that she is merely seeing a façade. It would then be true to state: “Ida knows that she is seeing a real building *rather than a mere façade*”. In a fake-barn case, however, the environmental conditions are such that Ida’s visual experience is not sufficiently discriminating to license such a knowledge ascription. Why so? Lewis would say that Ida is improperly ignoring certain possibilities; Schaffer would say that the salient contrast for the knowledge attribution is the one between fake and real barns. This implies that, on Schaffer’s account as well as on Lewis’, it is true to say, regarding the fake barns scenario, “Ida does *not* know that she is seeing a real barn rather than a mere façade”, because a salient, knowledge-threatening element has not been ruled out.

Finally, let us move to lottery cases. Imagine Jane, who has just purchased a lottery ticket, finds herself in the situation portrayed in (2). Her intuitive lack of knowledge is explained, in Lewis’ contextualism, by means of her improperly ignoring a possibility that is compatible with her evidence: namely, that the ticket will win. As Lewis explains (p. 557), either every one of these possibilities may be properly ignored, or none may. But one of them may not properly be ignored: the one that actually obtains. So given that actuality may not be properly ignored, neither can the possibilities that saliently resemble it.

Is Lewis’ suggestion to the effect that fake-barn cases and lottery cases are so closely linked with Gettier cases on the right track? In order to establish this, a careful comparison between Gettier and lottery- and fake barn-cases respectively needs to be conducted. This is what we set out to do in sections 3 and 4. Is Lewis’s contextualist diagnosis of Gettier cases viable to begin with? This question is tackled in sections 5, 6 and 7, where we point out some limitations of contextualism (including Lewis’ version) in capturing the ultimate source of Gettier problems.

3. Gettier and lotteries

In this section, we compare lottery cases with Gettier cases, arguing that the parallel put forward by Lewis only works up to a certain point. As it will emerge, the key difference lies in the *luck* component.

First of all, recall that in Lewis' view, when the subject states "I know my ticket will not win", her self-attribution is false because she is improperly ignoring the possibility that the ticket will win. Cohen (1988, p. 106) provides a similar explanation when he says that the statistical information concerning lotteries makes the possibility of error salient. DeRose (1996, p. 569) too voices a similar view when he says that we judge that the subject doesn't know because we realize she would have the same belief even if her ticket were in fact to win, thus being insensitive to the obtaining of certain possibilities.

An alternative way of spelling out the problem is offered by Hawthorne (2005) – who in turn draws on Vogel (1990). Key to this explanation is the so-called Parity Reasoning principle, which kicks off with the question: what special reasons does the subject have to believe that her ticket will lose? Apparently, none. Therefore, if she knows her ticket will lose, she has to know that any other ticket will lose, on pain of being arbitrary. Yet, it is unacceptable that the subject knows that every ticket will lose. Therefore, it should be inferred that the subject does not know her ticket will lose either.

Now, Parity Reasonings are quite sophisticated: they take some time and some effort to be spelled out. Philosophers may be able to create conversational contexts where the epistemic standards are suitably high as to demand that the Parity Reasoning be considered and dispelled; however, in less philosophical circumstances, it may be legitimate for us to deem sentences like "S knows that her ticket will lose" as true. This, Hawthorne observes (2005, p. 18), is especially likely to happen when certain practical aims are at stake. Suppose Sarah suggested to Ida, who is planning to buy a new car, to purchase a lottery ticket to get the cash she needs. This would seem like an insane advice. Ida would be well within her right to claim: "I know I am not going to win", at least relative to these practical purposes. Indeed Hawthorne suggests that, if one's plan or practical deliberation were to depend on a lottery win, then the huge probability that one's ticket will lose would suffice for one to *know* that one's ticket will indeed lose.

These considerations provide support to the view that contextualism can account for the problem underlying lottery cases. The contextualist is in a position to say that, (a) to the extent that the standards in the attribution context are suitably high, it is true to claim that the subject does not know she will lose – because (following Lewis, Cohen and DeRose) she cannot rule out certain salient possibilities or, alternatively (following Hawthorne and Vogel), she cannot avoid the Parity Reasoning; (b) conversely, when the standards in the attribution context are suitably low, it is true to state that the subject knows she will lose the lottery.

The full details of the contextualist account in play are not essential here. What is important is that the lottery problem can be accounted for with the resources typical of a contextualist theory, namely the idea of the attributor-sensitivity of the truth-value of knowledge ascriptions and the “shiftiness” of epistemic standards.

Can contextualism equally well articulate the Gettier problem? In Lewis’ view, it seems like it can. Lewis contends that the subject lacks knowledge of the gettierized proposition because, from the attributor’s standpoint, she is improperly ignoring a possibility that saliently resembles other possibilities that are not properly ignored. So far, Lewis’ reconstruction of the problem behind Gettier examples exactly parallels his reconstruction of the problem behind lottery cases. There is nothing wrong with this parallel. It seems that the subjects involved in Gettier cases *are* indeed guilty of the error Lewis imputes them. Lewis’ explanation is therefore so far accurate and we will not question it. The issue to be discussed is whether it is *exhaustive*. Gettier cases do not seem to just present a problem of improperly ignored possibilities. They are more than that. The commonly acknowledged source of the problem is that the subject knows by pure “luck”.

The notion of luck in epistemology has been receiving increasing attention in the last decade, allowing philosophers better to understand this phenomenon and its relation to knowledge.¹ As Duncan Pritchard (2005) notes, some forms of luck are epistemically benign, while others are malignant. For instance, it does not threaten knowledge that the subject’s capacity to acquire knowledge is the product of luck – as when a subject fortuitously survives dangerous circumstances, and as a result becomes capable of knowing certain facts (pp. 135-6). Neither does it threaten knowledge that the evidence is merely accidentally presented to the subject – as when somebody accidentally acquires evidence by overhearing a conversation (pp. 136-7). By contrast, it *does* threaten knowledge that a subject, who acquires

1 For an overview, see Rabinowitz (2011); Broncano-Berrocal (2016); Broncano-Berrocal and Carter (2017).

a belief based on a certain body of evidence, gets it right only as a matter of luck, in the sense that, although she is right, she could have easily been wrong. This is what Pritchard calls *veritic epistemic luck*, which he describes in modal terms as follows:

the agent's belief is true in the actual world, but [...] in a wide class of nearby possible worlds in which the relevant initial conditions are the same as in the actual world—and this will mean, in the basic case, that the agent at the very least forms the same belief in the same way as in the actual world [...]—the belief is false. (2005, p. 146; see also Pritchard 2015, p. 98)

We wish to argue that lottery cases and Gettier cases importantly differ with respect to the role of veritic epistemic luck. Consider again the stopped watch case. The subject believes that it is 8:29; her belief is justified on the basis of her consultation of the watch, a procedure that is in normal circumstances reliable; however, to put it in Pritchard's terms, it seems like in (almost) all close possible worlds where she acquires her belief as she actually does, namely through a glance at the watch, she wrongly believes that it is 8:29. This implies that her belief that it is 8:29 is actually true by sheer chance, simply because the actual world “plays in her favour” by instantiating a possibility that is not instantiated in the nearby worlds.

Think now of the lottery case. The subject believes her ticket will lose; she is justified on the basis of statistical information about the workings of lotteries; yet, the fact that she believes truly is not lucky at all: it is the most expected result! We may say, again in Pritchard's terminology, that in almost all the nearby possible worlds where she believes she will lose on the same grounds, she gets it right. Of course, the fact that her true belief is not lucky does not yet mean that she *knows* that her ticket will lose. It is compatible with what we have said so far that the subject has a non-lucky true belief that does not amount to knowledge, perhaps because a further requirement about strength of evidence is in place.

Interestingly, Pritchard contends that the belief of the lottery subject is in fact lucky, at least in one sense of the word. He first distinguishes between the probability of an event and its modal closeness, that is, how much of the actual situation needs to change in order for a certain possibility to obtain. Then he argues that, although losing the lottery is a very probable event, the truth of one's belief to the effect that one will lose is still the result of

luck, because *very little would have had to change* in order for the winning possibility to be instantiated – after all, it was only a matter of which configuration of balls would have been drawn (Pritchard 2005, p. 163; 2015, p. 97; forthcoming p. 3). Although this is an intriguing analysis, it seems to us that it can be trumped by the following consideration: there are very few nearby possible worlds where the subject wins. So, even though the winning obtains in virtue of a very tiny variation from one nearby world to the next, it is still true that the nearby modal space is almost entirely occupied by losing worlds. At the very least, the idea that only a tiny change would have sufficed to turn a losing ticket into a winning one, and the idea that there are many more worlds where the ticket loses than worlds where the ticket wins are in tension with one another, where this renders luck intuitions extremely unstable. The burden of proof seems to lie on Pritchard to explain why considerations of modal nearness should override considerations that have to do with a prevalence of losing alternatives in the surrounding modal space. As is formulated, then, the proposal is doubtful, for it is not yet clear that the subject who believes truly that she lost the lottery is in any way lucky.

The conclusion of this section is that luck is a central component of Gettier cases, which however seems (*pace* Pritchard) absent from lottery cases. In order to achieve a fulfilling explanation of Gettier examples, the contextualist should therefore strive to provide an account of this phenomenon by accommodating it in its theory. Can the phenomenon of epistemic veritic luck be satisfactorily captured in the contextualist framework? We postpone an answer to this question to sections 5, 6 and 7. In the next section, we compare Gettier cases with fake-barn scenarios.

4. Gettier and Fake Barns

The aim of this section is to compare Gettier cases and fake-barn cases. As it will emerge, the two scenarios share the luck component, with only minor and marginal differences. Given these similarities, it will be urged that, if it is possible to give a rendering of fake barn cases in contextualist terms – as some have proposed –, then it should be possible to do the same with Gettier cases.

It is easy to realize that both fake-barn cases and Gettier cases involve epistemic veritic luck, or believing the truth by accident. Applying Pritchard's formulation, we may see that in both cases, in (almost) all nearby possible worlds where the subject acquires the belief as she actually does – by seeing a barn façade or by glancing at the clock – she ends up with a

false belief. To be sure, the two examples differ in a number of respects: in the fake-barn case, the subject is supposedly exercising her visual capacities from a distance, where this increases the chances of error; in the stopped-clock case, the subject glances at the clock from a much less problematic distance. Moreover, the subject in the fake-barn scenario seems overall in a more epistemically hostile environment, given the extent and unpredictability of the deception he is exposed to; the stopped-clock case seems, in comparison, a rather mundane epistemic mishap. Still, it is conceivable that these differences be erased or significantly lessened, where this would reveal the strong structural similarity between the two cases.

In light of these similarities, it seems that, if fake barn cases can receive a contextualist treatment, the same should happen for Gettier cases. Greco illustrates how the verdict to the effect that the fake-barn subject does not know that there is a barn ahead of him may depend on assumptions about what is at stake:

[S]uppose that we are government employees, charged with counting barns in the area for the purposes of determining property taxes. Suppose also that barn facades are not taxed in the same way that working barns are. However, Henry is a new employee who does not realize that the area is populated with fake barns, and who has not yet received the special training needed to distinguish barns from barn facades. In this context, Henry sees a barn from a hundred yards and pulls out his log to record this. ... [I]f Henry were in this context to say, "I know there is a barn over there," we would not view his claim as true. But now consider a different practical reasoning context. Still in Fake Barn Country, we are working on a farm where we know that there are no barn facades. In fact, we know that there is only one structure on the property – a working barn. We and Henry are charged with getting a cow back to the barn. In this context, Henry sees a barn from a hundred yards and starts walking the cow in that direction. Now ... if Henry were in this context to say, "I know there is a barn over there," we most likely would view his claim as true. (Greco 2009, p. 80)

A discussion of the merits of this contextualist treatment would lead us too far astray, so we will not pursue it. What matters to us for present purposes is that, given the strong resemblances so far detected in fake barn cases and Gettier cases, the contextualist should be able to provide a similar story with, for instance, the stopped clock case or the Nogot-Havit

case. Specifically, the contextualist should be able to offer an account whereby knowledge denials in Gettier cases (as well as ascriptions, if any) can depend on features of the attributor's context – for instance, stakes, epistemic standards, properly ignored possibilities. As we have seen in section 2, Lewis' contextualism is certainly an option here. It is now time to assess this option.

5. Gettier and Contextualism: From Attributor-Sensitivity to Subject-Sensitivity

As a point of departure, let us first discuss an objection to contextualist accounts of Gettier-style cases that has been moved by Cohen (1998, p. 297): The contextualist theory says that the possibilities salient in the attributor's context determine whether or not the knowledge attribution is true. Thus, if in context C_1 the attributor is attending to the possibility that the clock tells 8:29 while it is really, say, 8:40, she will “call the shots” on the evaluation of the knowledge ascription. In C_1 , she would be right to say: “The subject does not know that it is 8:29”. However, suppose that in context C_2 the attributor was (appropriately) not attending to this very possibility. In the contextualist framework, it would follow that in C_2 , it would be true to state “The subject knows that it is 8:29”, despite the obtaining of the Gettier conditions.

This, Cohen maintains, “is strongly counterintuitive” (1998, p. 298), for it seems implausible that a subject might *know* – in any sense of the notion – that it is 8:29 when the conditions described above obtain, even if the attributor is not paying attention to some knowledge-threatening possibility.

It could be replied that the trouble for Contextualism is really just its exclusive focus on the ascriber, whose standards “call the shots” in the knowledge ascription. This is what primarily creates the problem in C_2 : on the one hand, in virtue of her being (appropriately) unaware of the Gettier possibility, the attributor truly ascribes knowledge to the subject; on the other hand, the subject's gettierized circumstances seem independently to suffice for denying knowledge. Contextualism should admit that, in this context, the subject's (improper) ignorings *trump* the attributor's (proper) ignorings. Thus, one could argue that if contextualism allowed to bracket the possibilities salient to the attributor and allowed the possibilities salient to the subject to affect the truth-value of the knowledge attribution, it would deliver the desired result in C_2 : namely, that the subject of the Gettier case does not know that it is 8:29. Cohen (1998) identifies the trouble of Lewis' contextualism with Gettier

cases as follows:

it is problematic to use a attributor-sensitive rule of relevance to solve the Gettier problem. There is no reason not to view the subject's failure to know in these Gettier cases as fixed across contexts of ascription—as holding regardless of who the attributor is. So any rule that solves the Gettier problem must be attributor-insensitive. (Cohen 1998, p. 304)

Cohen's suggestion is that Lewis' contextualism may better capture Gettier problems by becoming more subject-sensitive. One way for contextualism to achieve this aim is, according to Heller (1999, p. 125), to assume as part of the attributor's very own interests a concern for the *subject's capability* to distinguish between, say, a situation where it is 8:29 and one where it is 8:40; or to distinguish between a situation where Nogot owns a Ford and one where he owns no Ford. Once this concern for the subject is integrated in the attributor's own interests, it will be part of the attributor's job to carefully look into the possibilities that the subject is or not ignoring, whether or not they are salient in the attributor's context.

Whether or not Heller's suggestion helps, our persuasion is that making contextualism more subject-sensitive would be insufficient in order fully to capture the problematicity of Gettier cases, for reasons we explore in the next section.

6. Gettier and Contextualism: Failure to Capture the Luck Component

Can a more “subject-sensitive” version of contextualism capture the real problem with Gettier cases? In order to be able to answer this question, we need to establish what the “real problem” underlying these cases is. Recall our comparison between lottery and Gettier scenarios. There we moved the first steps towards making explicit at least one of the main problems of the Gettier examples, namely, *epistemic veritic luck*. We also noticed that epistemic veritic luck is shared by fake barn cases. There is therefore some pressure for contextualism to account for this element in its own framework.

Suppose we recruited the “subject-sensitive” version of contextualism suggested by Cohen and Heller in order to account for epistemic veritic luck. Would it provide a fulfilling

account? Our persuasion is that it would not: to explain Gettier cases exclusively in terms of improperly ignored possibilities would seem to offer an incomplete and unfair reconstruction of the case.

Our argument for this claim is quite simple: think of the elements that compose the stopped clock case. (1) The subject forms the belief that it is 8:29. (2) This belief is justified. (3) It is also true by accident. (4) The subject is ignoring the possibility that the clock says 8:29 and it is really 8:40. Now, suppose somebody asked “Why does the subject fail to know that it is 8:29?”. We might answer: “Because she is ignoring a possibility she should be considering (and ruling out)”. By giving this explanation, we would have taken care of the intuition of lack of knowledge, but we would have failed to capture what is peculiar of the Gettier situation. In effect, our explanation would seem to locate the stopped-clock case on the same level as, say, a bank case, where the attributor denies knowledge to the subject in virtue of the fact that the subject fails to consider, and therefore cannot rule out, the possibility that the bank will be closed on Friday despite evidence indicating that it will be open. But clearly, there is more to a Gettier case than the ignoring of certain possibilities (listed as element (4) above). There is also what we have listed as element number (3), namely, the accidental truth of the belief. This is exactly what sets a Gettier case apart from a bank case, and thus should be the target of an explanation aiming to account for the subject’s lack of knowledge. Therefore, explaining the Gettier subject’s lack of knowledge in terms of the ignoring of certain possibilities will have the cost of neglecting the most distinctive detail of the story, thus offering an incomplete and unfair representation of it.

The contextualist might retort that there is a systematic connection between element (3) and element (4), to be characterised as follows: the subject’s belief to the effect that it is 8:29 is accidentally true. This affects the epistemic standards in the subject’s scenario in such a way that she would have to consider and rule out certain possibilities which he would not have had to rule out absent the luck component. (The natural candidates here are those possibilities where the subject embraces a false belief despite using the same method of belief acquisition: for instance, the possibility whereby it is 8:40 but she still comes to believe that it is 8:29 based on her consultation of the clock.) Since these possibilities remain improperly ignored, the subject fails to know. In this scenario, luck acts as an element capable of modifying the epistemic standards operative in the subject’s situation, and consequently capable of affecting which possibilities are properly ignored and which are not. Yet, even if luck could play this role, this would render the changing of the standards a quite explanatorily

idle step. For the fact that the subject believes the truth by luck seems already independently sufficient to deny knowledge, with no need to further process this element in the contextualist machinery of “shifty standards” and “improperly ignored possibilities”. That is to say, the fact that the subject’s belief is true by chance seems already effectively to explain why she fails to know that it is 8:29, with no need to mention a change in epistemic standards and the consequent improper ignoring of certain possibilities.

The conclusion of this section is that either a contextualist explanation of Gettier cases is incomplete and fails to do justice to the distinctive feature of these examples – namely veritic epistemic luck; or, if it attempts to incorporate the luck element into the explanation, it turns the typically contextualist account in terms of shifty standards (and improper ignoring of possibilities) into an idle wheel. This implies that contextualism offers an unsatisfactory account of all the cases reviewed so far where luck bears on the knowledge attribution: stopped clock cases, Nogot and Havit cases and fake barn cases.

Having achieved this result, our discussion could potentially stop here. Yet, there is a further aspect to Gettier cases that has not emerged so far, and that reveals a further way in which contextualism has a hard time accounting for Gettier cases.

7. Gettier and Contextualism: Failure to Capture Lack of Causal Connection

So far, we have argued that contextualism inadequately accounts for the luck component of Gettier cases. Based on our discussion so far, Gettier cases comprise stopped-clock examples, Nogot and Havit cases as well as fake barn scenarios. This is not completely accurate, though. There is an important, further difference that has not emerged so far between fake barn cases and the rest of the “standard” Gettier examples. We contend that contextualism fails to account for this differentiating feature. This provides additional reasons to think that contextualism’s relation with Gettier cases is highly problematic.

The further difference we wish to emphasise lies in the *connection* between the subject’s evidence and her true belief. Arguably, while this connection non-problematically obtains in the fake barn case, it fails to obtain in Gettier examples.

Consider again the subject in the fake barn country: he looks from a distance and happens to spot the only real barn in the surroundings. Now, if we ignore possible defeaters and set aside epistemic veritic luck, we notice that all goes well when it comes to the *link*

between the subject's evidence and the truth of his belief. The subject acquired some evidence as to the presence of a barn when there indeed is a barn ahead of him; so the evidence in this case is *truth-conducive*. Thus, we can say that the subject acquires the true belief that there is a barn ahead of him "thanks to" the fact that he had that particular evidence. Both the evidence and the process of its acquisition were connected to the true belief in an acceptable way. The fake barn subject is therefore *creditable* to his true belief. Notice that this is likely to generate a clash of verdicts about fake barn examples: on the one hand, the subject being lucky would seem to push in the direction of a knowledge denial; on the other hand, the subject's creditability would seem to push in the direction of a knowledge ascription. If this is so, then we are happy to leave it open whether fake-barn cases are cases of knowledge; our consideration from now on will be restricted to more "traditional" Gettier cases, like the stopped-clock or the Nogot-Havit case.

The link between evidence and true belief that we detected in fake barn cases is absent in the other Gettier cases we have been scrutinising. John Greco brings this point out in his own terminology:

in Gettier cases, S believes the truth, and S believes from an ability, but S does not believe the truth because S believes from an ability. [...] the "because" is here intended to mark a causal explanation. (Greco 2009, p. 74).

Thus, in Gettier cases, the subject believes the truth on the basis of evidence acquired through an epistemically acceptable method, but she does not come to believe the truth *because of*, or *thanks to*, the evidence acquired through that method. In the stopped-clock case, the subject comes to believe that it is 8:29 based on her looking at the clock. Her belief is true, her method of belief acquisition is generally reliable and her evidence is sufficient; however, that very evidence and the process through which she acquired it "deserve no credit for" causing her to believe something true. For hadn't it been for a lucky coincidence (it actually being 8:29), the subject's evidence would have led her to a false belief. Similarly, in the Nogot and Havit case, the subject forms the belief that either Nogot or Havit owns a Ford, based on evidence she has for the first disjunct of the proposition, consisting of visual experience and memories of Nogot driving a Ford. Although this evidence decently supports the first disjunct, the disjunction is true because the second disjunct is true. Here too, the evidence the subject

has “deserves no credit” for leading her to believe the truth, because it provides epistemic support for the false disjunct and the fact that the other disjunct (and thus the whole disjunction) was true was an altogether fortuitous result.

Now, it seems to us that this defective liaison between evidence acquisition and true belief acquisition is simply not articulable in the contextualist framework. It is one thing for the subject to have or lack the quality and amount of evidence demanded by the standards relevant in a certain context (whether the subject’s or the attributor’s); it is another thing to have evidence that defectively connects with the true belief she acquires on its basis. The master consideration supporting the latter claim is the following: It is perfectly possible to imagine that in a Gettier case the subject has exactly the quality and amount of evidence demanded by the relevant contextual standards – which may range from ordinary to extremely low –, but the evidence is defectively linked to the true belief, thus endangering knowledge for independent reasons. These seem precisely the conditions in which the problematicity of the Gettier case best emerges. In the stopped-clock case it indeed seems that, at least relative to *ordinary* epistemic standards in the context of the ascription, the subject has the adequate quality and amount of evidence required for her to know that it is 8:29; and the same goes for the Nogot and Havit case. The subject’s evidence is unassailable given ordinary standards; what engenders the intuition of a lack of knowledge is something *other* than the failure to meet epistemic standards: plausibly, the defective nexus between evidence and (lucky) true belief.

The contextualist might reply that, when a subject’s evidence fails to be the appropriate cause for the subject’s true belief, the evidence fails certain contextual standards required for belief to count as knowledge. Thus, this defective connection between evidence and true belief has a correlate in terms of epistemic standards within the contextualist framework. Yet, it is difficult to see the epistemic standards in any Gettier scenario as being significantly higher than they are in ordinary circumstances. This already casts suspicion on the idea that, with Gettier cases, the theorist is managing to raise the contextual standards in a way comparable with, say, bank examples. Moreover, it is extremely difficult to imagine a context where these standards are loosened: there would hardly seem to be any situation in which the epistemic standards could be *so relaxed* as to allow that, even if evidence and true belief bore such a defective relation, the subject would still have knowledge. That evidence and true belief are related in a specific, non-Gettier-like way does not seem to be a requirement that may legitimately hold in some contexts and not in others, as one would

expect to happen in a contextualist theory. Rather, this requirement would seem to hold in *all* contexts of ascription, no matter how low their standards, thus violating the “shiftiness” that is typical of context-sensitive aspects of knowledge.²

We therefore conclude that, if what sets Gettier cases apart from fake barn cases is the defective connection between evidence and true belief, then the fact that this defective connection cannot satisfactorily be captured by a contextualist account gives us a further reason to believe that contextualism is not well positioned to offer a complete explanation of the problem underlying “traditional” Gettier cases – like stopped-clock cases and Nogot and Havit cases.

Perhaps, though, we have not considered one last possible type of contextualism, which could stem from a *rejection* of the distinctive Gettier intuition to the effect that the subject lacks knowledge. Stephen Hetherington (1999, 2006) has argued that, in Gettier scenarios, although it is false to say that S knows that *p* in the standard sense of knowing, it is true to say that S knows that P in a *non-standard* sense of knowing. Such a non-standard sense of knowing would be one explicitly admitting “lucky knowledge”, that is, knowledge that could have easily failed to be so.

The crucial point for the present discussion is whether this position can be developed into a form of contextualism. One way to do so would be to argue that “know” is *ambiguous*, and that depending on the context, either the standard sense of “know” may be used, which demands an absence of luck, or a less strict sense of “know” may be employed, which does not demand an absence of luck. In the spirit of contextualism, this thesis would require further backing by means of evidence stemming from a close examination of our linguistic practices. It would also require a significant departure from the overall monistic attitude concerning knowledge contextualists have had so far. For they have usually maintained that “knowledge” has a univocal sense and that it picks out just one kind of epistemic property, which can then

2 Contextual variation need not have to do with epistemic standards; importantly, it could concern causal relations. Greco (2004, pp. 391-2) defends a contextualism about “because” whereby, depending on the context, the subject might or might not have knowledge based on whether her evidence acquisition appropriately caused her true belief – typically, through the exercise of intellectual virtue. Yet, even this contextualism would not help with Gettier cases, for in order to do some explanatory work, it would have to be compatible with a context in which “because” is applied in such a way that evidence acquisition for the Gettier subject is appropriately causally connected to the true belief. Yet, there seems to be no context where the conditions of application of “because” are so relaxed as to admit that evidence acquisition in a Gettier case appropriately causally connects with true belief, for this would imply excessively disregarding the subject’s intellectual abilities and virtues. If there is no such context, then it seems that the same explanation available to a contextualist is also available to someone who is not a contextualist about “because”.

be either attributed or denied depending on contextual shifts of standards.

Another possibility would then be to argue that “know” has only *one* sense and picks out just one property. Still, its sense is, upon reflection, more permissive than one might have thought. This would imply that subjects in Gettier scenarios know that *p* in the same sense of “know” in which subjects in non-Gettier scenarios know that *p*, in virtue of the fact that the only sense of “know” available is, upon philosophical reflection, sufficiently broad as to encompass cases in which knowledge obtains by luck. Hetherington’s own proposal might be read exactly in this fashion, when he remarks: “There is a real possibility of our having been thinking about knowledge in a way that is more conceptually restricted than we have realised” (2006, p. 219; see also p. 223; see also 1999, p. 586). Since Plato’s rejection of Theaetetus’ second definition of knowledge, philosophers have been quite skeptical of such a possibility. Hence, its proper defense would require substantial elaboration.

An examination of the latter two proposals falls outside the scope of the present paper, whose goal was to examine what extant contextualist theories have to offer in the way of a diagnosis of Gettier cases.

Conclusion

We have explored the question whether contextualism can provide a satisfactory account of the Gettier problem. Our point of departure was Lewis’ contention that Gettier cases crucially involve the improper ignoring of certain uneliminated possibilities, where this feature also makes lottery and fake-barn scenarios very close to Gettier cases. By comparing Gettier cases with lottery cases and with fake-barn cases respectively, we have isolated some features that are distinctive of a gettierized scenario, and that should be accommodated by a contextualist account. These features are (i) epistemic veritic luck and (ii) a defective relation between evidence and true belief. We have found no convincing way in which contextualism could accommodate these features within its framework; we have therefore concluded that a contextualist account of Gettier cases would at best be partial or incomplete.

References

Broncano-Berrocal (2016). Luck. Internet Encyclopedia of Philosophy.

- Broncano-Berrocal & J. Adam Carter (2017). Epistemic Luck. In Tim Crane (ed.), Routledge Encyclopedia of Philosophy.
- Cohen, S. (1988) How to Be a Fallibilist, *Philosophical Perspectives* 2 (1988) pp. 91-123.
- Cohen, S., (1998), “Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery”, *Australasian Journal of Philosophy*, 76(2): 289–306.
- Cohen, S., (1999), “Contextualism, Skepticism, and The Structure of Reasons”, *Philosophical Perspectives* 13: Epistemology: 57–89.
- DeRose, K. (1992), “Contextualism and Knowledge Attributions”, *Philosophy and Phenomenological Research*, 52(4): 913–929.
- DeRose, K. (1995), “Solving the Skeptical Problem”, *The Philosophical Review*, 104(1), 1–52.
- DeRose, K. (1996) Knowledge, assertion and lotteries, *Australasian Journal of Philosophy*, 74:4, 568-580.
- DeRose, K. (1999), “Contextualism: An Explanation and Defense,” in *The Blackwell Guide to Epistemology*, J. Greco and E. Sosa, eds., Malden MA, pp. 185–203.
- Dretske, F. (1981) *Knowledge and the Flow of Information*, Oxford: Blackwell.
- Gettier, E. L. (1963). “Is Justified True Belief Knowledge?” *Analysis* 23: 121-3.
- Goldman, A., (1976), “Discrimination and Perceptual Knowledge”, *The Journal of Philosophy*, 73: 771–791. Reprinted in Goldman 1992, pp. 85–103.
- Greco, J. (2004) A Different Sort of Contextualism. *Erkenntnis* 61 (2-3): 383-400.
- Greco, J. (2009), *Achieving Knowledge*, Cambridge and New York: Cambridge University Press.
- Hawthorne, J. (2005) *Knowledge and Lotteries*. New York: Oxford University Press.
- Heller, M. (1999), “The Proper Role for Contextualism in an Anti-Luck Epistemology”, *Philosophical Perspectives*, 13 (Epistemology): 115–129.
- Hetherington, S. (1999) Knowing failably. *The Journal of Philosophy* 96: 565–587.
- Hetherington, S. (2006) Knowledge that Works: A Tale of Two Conceptual Models, in

- Hetherington, S. (ed.) *Aspects of Knowing. Epistemological Essays*. Oxford: Elsevier.
- Lewis, D.K. (1996), "Elusive Knowledge", *Australasian Journal of Philosophy*, 74(4): 549–567.
- Neta, R. (2003), "Contextualism and the Problem of the External World", *Philosophy and Phenomenological Research*, 66(1): 1–31.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2015) Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172:93–111.
- Pritchard, D. (forthcoming) Epistemic Risk. *The Journal of Philosophy*. Page references from penultimate draft.
- Rabinowitz, D. (2011). The Safety Condition for Knowledge. *Internet Encyclopedia of Philosophy*.
- Russell, B. (1948) *Conditions of Knowledge*. Chicago: Scott, Foresman.
- Schaffer, J. (2004). From Contextualism to Contrastivism. *Philosophical Studies* 119 (1-2):73-104.
- Schaffer, J. (2005). Contrastive Knowledge. In Tamar Szabo Gendler & John Hawthorne (eds.), *Oxford Studies in Epistemology* 1. Oxford University Press. pp. 235.
- Vogel, J. (1990) 'Are there Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer).

Acknowledgments

Thanks to Giorgio Volpe, Chris Kelp and Duncan Pritchard for their helpful comments on earlier drafts of this paper. Delia Belleri gratefully acknowledges support from the Humboldt Foundation.