Accidentality and Knowledge After Gettier

Leandro De Brasi

UDC: 165.18

Abstract

Fifty years on since the publication of Gettier's famous article and there is still no widespread agreement among philosophers on some set of conditions that block Gettier cases and the like. In fact, there is much pessimism about some such project. The aim of this paper is to show that this pessimism is misplaced. I argue that an account that posits a reflective endorsement of the knowledge-yielding procedures that can be met at the social level can capture, what I here call, the *non-accidentality desideratum*. The lesson we can draw is that the project of adding some accidentality-blocker to true belief to have knowledge is not doomed if we give up the individualistic approach of much mainstream epistemology.

Keywords: knowledge, epistemic luck, Edmund Gettier, reflective endorsement, global reliabilism

Epistemology has experienced a bit of a renaissance within the last few decades and much of this can be traced back to Edmund Gettier's famous paper, "Is Justified True Belief Knowledge?". Suddenly there was a puzzle that became the centre of attention; namely, the nature of knowledge, and there has been no shortage of proposals attempting to identify what has to be added to true belief in order to get knowledge. Different proposals put forward different "no accident" clauses that eliminate the knowledge-threatening epistemic luck that problematic cases, such as Gettier's own, exploit.

Fifty years on since the publication of Gettier's article and there is still no widespread agreement on some set of conditions that block these problematic cases. In fact, there is much pessimism about such project since the proposals have their own counter-examples with no sign of developing a successful one (e.g., Millar; Williamson). Moreover, the proposals are anyway likely to be quite complex and gerrymandered (and the more successful they are at dealing with counter-examples, the more complex and gerrymandered they are likely to be), making it hard to understand why we would have a concept that referred to such a thing (Hyman), let alone care about it (Kvanvig, *The Value*).

The aim of this paper is to show that this pessimism is misplaced. The project of adding some accidentality-blocker to true belief to be in a position to have knowledge

is not doomed if we give up the individualistic approach of much mainstream epistemology. The social approach here put forward, I suggest, allows us to successfully address the *non-accidentality desideratum*, as I shall call it. Thus, I claim that there is an important lesson to be drawn from this. We can make progress understanding the nature of knowledge by not neglecting the realities of social interaction.

The paper proceeds as follows. In §1, I introduce the non-accidentality desideratum. In §2, I present the Gettier cases and the type of knowledge-threatening epistemic luck that we should rule out. In §3, I put forward an account of knowledge that exploits social elements. In §4, I show how the proposed account meets the non-accidentality desideratum by showing how it deals with distinct sorts of cases. In §5, I offer some brief concluding remarks.

1 Non-Accidentality

Knowledge is ordinarily thought to be non-accidental, i.e., when we know, we don't have an accidental grasp of the truth. Crudely put, knowledge is incompatible with luck or chance, such as being the result of some lucky process.[1]

Most philosophers would accept this as a *desideratum*. Indeed, some have even provided analyses of knowledge that exploit this feature, such as: S knows that *p* iff it isn't at all accidental that S is right that *p* (Unger 158). But unfortunately no such account will do. The problem is that, even if they are accurate, they are neither theoretically illuminating nor practically useful (Zagzebski, "What is Knowledge?"). So, as Bernard Williams says, we "offer the "no accident" clause not as part of an analysis but [...] as a label for a class of conditions, the general requirements on which need to be spelled out with greater precision" (7). This is exactly what needs to be done—provide an anti-luck account of knowledge and elucidate the vague and not illuminating "non-accidental."

Nevertheless, although we commonsensically think that the end-result isn't knowledge if there is a fortuitous connection to the truth, we take some forms of luck to be benign, as when one is lucky to come across evidence that *p* (Unger). Still, our account should be in the business of dealing with Gettier-type cases since they seem to involve knowledge-threatening luck: in a sense the Gettiered subject gets things right by accident.

2 Veritic Luck

Gettier offers a couple of cases where one can have a justified true belief but intuitively one doesn't seem to know (122-3). Here is one of Gettier's cases: Smith is justified in

[1] I shall use 'accidentality' and 'luck' interchangeably.

believing that Jones will get the job and that Jones has ten coins in his pocket. Smith then forms the justified belief that the person who will get the job has ten coins in his pocket. But Jones doesn't get it, Smith does, who happens to have ten coins in his pocket. So Smith has a justified true belief but, significantly, he doesn't seem to know.

His cases, Gettier claims, refute the Justified True Belief analysis, since justification, belief and truth aren't sufficient for knowledge. Now, Gettier's cases exhibit two general features: fallibility and accidentality. The first feature allows a belief to have positive epistemic status even though not all error-possibilities have been eliminated and so allows for the possibility that one can reach a false belief by exploiting a legitimate procedure (Gettier 121). The second feature allows a belief to be true by luck.[2] So, in Gettier's cases, one seems to achieve knowledge-constituting status about a belief that only by accident happens to be true.

Gettier cases then share both these features. When I form the belief that a sheep is in the field, I do so on the basis of, say, a reliable visual procedure that is meant to confer the relevant epistemic status but, unbeknownst to me, what I am actually looking at is a dog that resembles a sheep. Luckily, however, there is a sheep in the field, and so my belief is true and possesses the epistemic status to be knowledge but importantly it seems that it isn't an item of knowledge (Chisholm). Therefore, we can also see why some characterize Gettier cases as double-luck cases (Zagzebski, "The Inescapability", *Virtues*): where some bad luck (this time the non-perfectly reliable procedure didn't deliver the goods, although it normally does) is cancelled by some good luck (it just happens to be the case that there is a sheep in the field).

Now it seems that if we want our theory to be Gettier-proof, then we would do well in eliminating either of those features.[3] However, it isn't attractive for most of us to eliminate fallibilism (cf. Sturgeon). Scepticism is as unwelcomed as accidentality (if not more). So, by far, the most attractive option to deal with Gettier cases seems to be to eliminate accidentality. Of course this can be and has been attempted in different ways. Here I propose an account that can handle accidentality that I have developed elsewhere. But before I introduce the account, some distinctions with regard to the different types of luck that might be in play are called for. This is important because not all kinds are knowledge-undermining (Unger).

Following Duncan Pritchard (*Epistemic Luck* 133ff.), we can differentiate at least four *benign* varieties of epistemic luck: 1) *content luck* (when "it is lucky that the proposition is true," i.e., that a certain fact occurs), 2) *capacity luck* (when "it is lucky that the agent is capable of knowledge," i.e., that she exists or has the pertinent abilities), 3) *evidential luck* (when "it is lucky that the agent acquires the evidence that she has

[2] See e.g. Dancy; Pritchard, *Epistemic Luck*; Unger; Zagzebski, "The Inescapability."

[3] Zagzebski ("The Inescapability", *Virtues*; cf. *On Epistemology*) seems to suggest that only infallibilist accounts can avoid these counterexamples; but see Howard-Snyder, Howard-Snyder, Feit.

in favour of her belief," i.e., that certain evidence becomes available) and 4) *doxastic luck* (when "it is lucky that the agent believes the proposition," i.e., that she forms the belief). Luck in these cases doesn't seem to threaten knowledge. But there is at least one type of epistemic luck that is knowledge-threatening: namely, *veritic luck*, which is "a matter of luck that the agent's belief is true."[4]

This is the knowledge-undermining luck that seems at play in Gettier cases (as the above "double-luck" description of the cases suggests), although, as we shall see below, there are different sub-types that seem incompatible with knowledge. Anyway, this is the kind of luck that we would need to eliminate if the account isn't to be subject to that sort of Gettier cases. So, given these cases which exploit veritic luck, we need to consider whether the account here proposed, and to be introduced next, can escape this kind of knowledge-threatening accidentality: whether knowers, according to our account, can't be veritically lucky believers.

3 Knowledge and Reflective Endorsement

It is often thought that a satisfactory account of knowledge should combine objective and subjective standards of appropriateness for a true belief to be knowledge (e.g., Fogelin; Russell). Normally this is understood in terms of a combination of *de facto* reliability and epistemic responsibility (e.g. Greco, *Putting Skeptics*; Zagzebski, *Virtues*). This combination seems desirable because, on the one hand, views that take into account only the subjective (or perspectival) aspect seem to fail to be appropriately normative. If Sid responsibly decides to use Tea-Leaf-Reading practices to find out the truth, those unreliable procedures do not seem to provide the knowledge-relevant normative status (Goldman, *Epistemology*; Greco, *Achieving*). On the other hand, views that take into account only the objective aspect are vulnerable to cases in which irresponsibility seems to rule out knowledge regardless of the reliability of one's belief-forming procedure (Zagzebski, *Virtues*). The well-known cases concerning reliable but epistemically naïve subjects, such as Norman (a reliable but naïve clairvoyant) and Mr. Truetemp (a reliable but naïve temperature-teller), suggest that reliability's positive epistemic status is swamped by irresponsibility due to their naivety with regard to the reliability of the procedures exploited (BonJour; Lehrer).

But this "standard way of looking at things" (Grimm 90) quickly gets us into trouble. Cases the pure reliabilist is quite fond of suggest this. For example, proprioception is a reliable faculty that many know nothing about it (not even about its existence), but we are nonetheless willing to attribute knowledge to those subjects when exploiting it. However, given reliability is necessary but not sufficient for the knowledge-relevant

[4] As Pritchard says: luck that "'intervenes' between the agent and the fact" ("Knowledge and Understanding" 36).

normative status of a true belief to be knowledge and that cases of reliable but epistemically naïve subjects motivate a responsibilist condition, we still seem to need some sort of admixture of reliability and responsibility. The problem, however, is to find a satisfactory combination of reliability and responsibility (M. Williams). But such combination is to be found by capitalizing on the real and ubiquitous human phenomenon that is the social dispersal of epistemic labour through time (De Brasi, "Reliability").

We cannot, as BonJour does, understand epistemic responsibility as demanding *the knower herself* to reasonably take the procedure exploited to be reliable, at the expense of over-intellectualizing knowledge.[5] After all, most ordinary subjects would often fail to satisfy this condition (even if allowed to satisfy it tacitly–BonJour 50). It seems that clear knowledge needn't be a reflective success of the subject: "Knowledge is a matter of responsiveness to the way the world is" (Roush 122), but not necessarily to reasons of the subject.[6] Having said that, we needn't give up BonJour's idea that we require some reflective endorsement of the knowledge-yielding procedures.

Although the personal reflective endorsement of procedures is implausibly demanding, its social analogue isn't. The proposed account then suggests that such reflective endorsement needn't be performed by every knower of the epistemic community. Instead we can all exploit, as we often do, different procedures rooted in the epistemic community, for which we lack positive reasons, as long as the reflective endorsement is met at the social level. That is, although the knower needn't possess the positive grounds for the endorsement of the procedures that she and other members of the community rely on, some member of the community does. These procedures are, I shall say, *socially endorsed*: in the sense that some subject or, more likely, a group of subjects of the knower's epistemic community have undertaken through time the positive epistemic work to reasonably take the community's established procedures to be reliable.

The clearest example of subjects who appreciate the rationale behind our belief-forming procedures and who would also promote their revision through time if regarded necessary, is that of regulative or ameliorative epistemologists, whose job is to actively engage in the project to remedy the deficiencies of our epistemic practices (to increase their truth-conduciveness). In particular, the social endorsement is sometimes the product of some sort of epistemic policing that prompts the correction or perfection of inadequate procedures. And this social endorsement means that one, as a member of the community, is permitted to exploit any of its rooted knowledge-yielding procedures even in the absence of personal positive reasons for their endorsement.

[5] "Part of one's epistemic duty is to reflect critically upon one's beliefs, and such critical reflection precludes believing things to which one has, to one's knowledge, no reliable means of epistemic access" (BonJour 42). In particular, the knower should at least reasonably judge the belief-forming procedures to be reliable from within her system of beliefs (BonJour 50, 123).

[6] See e.g. Ayer; Goldman, *Epistemology*; Lewis; Millar.

This anti-individualist approach to knowledge-relevant responsibility yields the following account: (as a first approximation and ignoring defeaters and factivity) S knows that *p* iff (i) S's belief that *p* was formed by a reliable procedure, AND EITHER (ii.a) the procedure is reflectively endorsed by some member(s) of S's community and (ii.b) reliance on the procedure is a standard or established practice of the community, OR (iii) S herself reflectively endorses the procedure.

This account exploits a division of epistemic labour where some members of the community responsibly endorse the procedures for others and so corrects some unfortunate individualistic simplifications of much current mainstream epistemology that are inapt for theorizing about knowers who are members of social communities. Significantly, given we are to satisfy some perspectival dimension on knowledge, allowing the endorsement condition to be met at the social level makes the burden of epistemic responsibility less hard to endure, and so we end up with a plausible account that can accommodate the relevant intuitions regarding knowing and non-knowing epistemically naïve subjects, such as those who exploit proprioception and clairvoyance, respectively (De Brasi, "Reliability"). Moreover, the account can be independently motivated (De Brasi, "Testimony") and importantly for our purposes it can, as we shall see below, handle accidentality. The details of the account aren't relevant here, since all we need to capture the non-accidentality of knowledge is an account that posits belief-forming procedures that are endorsed to fulfil a certain job.

4 Gettier Cases

So, what can this account say about Gettier cases? In order to answer the question we need to note that the knowledge-yielding procedures are endorsed to connect us to the fact that makes the target proposition true. This is because the procedures are meant to be truth-discriminating—that is, they discriminate truths from falsehoods.[7] This notion of discrimination is central to reliability (Goldman, "Discrimination," *Epistemology*). As Alvin Goldman says, "To be reliable, a cognitive mechanism must enable a person to discriminate or differentiate between incompatible states of affairs" ("Discrimination" 771). So this should be seen as a refinement on our understanding of knowledge-yielding procedures: a procedure for *p* can discriminate *p* from other state of affairs where *p* is false.[8]

[7] Indeed, if this capacity were lacking, one would be achieving the truth accidentally and not have knowledge (McGinn).

[8] Consequently, a "knowledge attribution imputes to someone the discrimination of a given state of affairs from possible alternatives" (Goldman, "Discrimination" 772).

Indeed, as mentioned, knowledge requires responsiveness to the world and discrimination is the natural option that allows us to achieve such responsiveness.[9] Knowledge requires a proper connection to the fact that makes the target proposition true, where a "proper connection" is a connection that allows us to discriminate such proposition from other state of affairs where it is false. So knowledge-yielding procedures provide us with the capacity to distinguish competing state of affairs and it is with this constraint in mind that the procedures are endorsed.

But Gettier cases are *ex-hypothesi* cases, in which no such connection occurs because they are cases of "double-luck." In other words, in these cases, the procedure fails to connect us to the fact (say, the sheep in the field) that makes the target proposition true. Now, given fallibilism, Gettier cases don't give us a reason for not endorsing the procedures we happen to exploit in those cases unless these cases represent (given our worldview) significant error-possibilities. And I take it they don't, so these procedures needn't be given up.

Having said that, in these cases, the procedures do fail to connect us to the relevant facts and so they don't work as they should. Just like an artefact can fail, in certain cases, to work in accordance with the way it was designed, so can the procedures. And Gettier cases are cases in which this is so. The Gettiered subject doesn't know since the procedure employed fails to work as it was conceived to do by not connecting her to the fact that makes the target proposition true. That is, in cases of double-luck, the subjects employ procedures that are endorsed to fulfil a certain function but fail to do so. In these cases, we can say that the procedures fail to work as conceived when endorsed.

4.1 Goldman-Ginet Cases

But there are other cases, which one might at first think are Gettier-*like*, that don't seem in fact to have the same structure as Gettier cases. One such case is the fake-barn case, where although one believes that it is a barn after seeing the only real barn among many fake ones, one doesn't seem to know (Goldman, "Discrimination").[10] So, unlike the above sheep case where the subject isn't looking at a sheep but the belief she forms ends up anyway being true, in the fake-barn case the subject does see the real barn about which she forms her belief. And the problem in the latter case seems to rest with the nature of the environment in which one comes to believe that it is a barn. Indeed, it is important

[9] Sensitivity, which roughly requires that if *p* were false, S wouldn't believe that *p*, is another way of being responsive to the world. But, Roush argues, safety, which roughly requires that if S believed that *p*, then p would be true, fails to be so since "it gets the direction of fit wrong for what knowledge is" (121). Regardless of this, neither of them seems the natural (or right) option; see below.

[10] These aren't cases where the subject lacks knowledge because of some defeater. The subject is unaware of the situation and we can stipulate that there is no reason why she should be. If this weren't so, the case would be explained by such defeater.

to distinguish at least two different types of problematic cases.[11] Firstly, those cases in which what seems to be exploited is some kind of *veritic double-luck*—a luck that "'intervenes' between the agent and the fact." Gettier's own cases and Chisholm's sheep case are of this kind, so we can continue to call them Gettier cases. Secondly, those cases in which what seems to be exploited is some kind of *environmental veritic luck*—a luck that concerns the environment in which the procedure generates the success. The paradigmatic example being the fake-barn case, so we can call them *Goldman-Ginet cases* (given Carl Ginet's involvement in Goldman's fake-barn case).

Consequently, what can our account say about Goldman-Ginet cases? Again, in order to answer the question we need to note that the knowledge-yielding procedures are responsibly endorsed to promote the truth given our worldview. Now this endorsement of procedures presupposes that certain background conditions are in place (say, that it isn't habitual for countrymen to build fake-barns) and, given these conditions, the procedures are reasonably taken to be truth-conducive. Indeed, these are the conditions that one naturally presupposes, given our shared worldview, to obtain when employing a procedure. And it is these background conditions presupposed by the procedures, i.e., the conditions for which the procedure was endorsed, that allow us to handle the denials of knowledge in Goldman-Ginet cases. After all, in these cases it seems the procedures would be exploited in situations that they weren't endorsed for, and so we can explain why knowledge isn't to be possessed.

The procedures are endorsed to be exploited *in certain environmental conditions* in which they are reasonably taken to be truth-conducive (hence allowing them to be legitimate[12]), but in Goldman-Ginet cases those environmental conditions are different if the subject is at the time exploiting a procedure that for *her* community (given their worldview and environment) is legitimate.[13] The subject then doesn't know since there is a mismatch of environmental conditions—a mismatch between the conditions for which the procedure was endorsed (say, a fake-barn free environment) and the conditions in which it is used (say, an environment that is replete with fakes). So the procedure to determine from a distance whether something is a barn is endorsed, given our worldview, with a certain type of environment in mind, which isn't, as the case states, one where there is an abundance of fake-barns (after all, attempting to individuate barns in such a fake-barn infected environment wouldn't be reasonably regarded as truth-conducive), and so we don't know in the fake-barn case.

[11] See e.g. Greco, *Achieving*; Pritchard, *Knowledge*, "Apt Performance," "Knowledge and Understanding."

[12] Of course, as mentioned above, in order to be legitimate procedures, they also need to be, as a matter of fact, truth-conducive.

[13] If the subject weren't exploiting a *legitimate* procedure of her community in this sort of case, then we could easily explain her lack of knowledge in the fake-barn case since merely looking at a barn from a distance wouldn't put us on a position to know that there is a barn.

So the account can also capture our intuitions concerning the Goldman-Ginet cases. But one might think that ours is a pseudo-solution since it still can't deal with the real problem–that even if we aren't in an epistemically inhospitable environment, we *could have been* in some such unfortunate environment (Pritchard, "Knowledge and Understanding"). It is just a matter of luck that we aren't, and this luck is meant to eliminate knowledge. Now, if this were so, our account would be in trouble. After all, if the procedure is being applied in the environment for which it was endorsed, and even if it is just a matter of luck that we happen to be in such an environment (and barring defeaters), then such application of the procedure won't eliminate knowledge. Fortunately, our intuitions don't seem to clearly follow that pattern.

Let us adapt the fake-barn case so that the subject is in Real-Barn County, where there are no fake-barns, and sees a real barn but she would, unbeknownst to her, instead be in Fake-Barn County if she had turned left at the junction 10 miles ago (as she could have easily done, since she is just driving around). My intuition is that she knows. And, although I don't expect everyone to share it, I do expect many to do so, since this just seems to be a case of benign luck: particularly, a case of capacity luck, when it is lucky that the agent is capable of knowledge by having the pertinent abilities.[14] Our strong intuitions about the subject in the original fake-barn case lacking knowledge are the result of *actually* exploiting the ability or procedure in the wrong environment. And this then is a case of "environmental luck" because, although the environment in which the procedure is being applied isn't the one for which it was endorsed, the procedure stills delivers in this instance the truth. Such environmental luck is knowledge-precluding and our account can do justice to this.

4.2 Harman Cases

However, Goldman-Ginet cases aren't alike another kind of cases, which we might call, *Harman cases*, such as the political assassination case, Donald's letters case, and Tom

[14] Cf. Sosa, *A Virtue*. In his account of (animal) knowledge as apt true belief, a belief is apt if it is accurate because it is competent or "adroit." What matters is the ability of the epistemic agent and that this cognitive ability is responsible for the cognitive success. A belief isn't apt if it isn't successful "sufficiently" because of the subject's adroitness (Sosa, *A Virtue* 79). It seems this account won't fully capture our intuitions about Goldman-Ginet cases, since environmental luck doesn't seem to interfere with the aptness of a performance (Pritchard, "Apt Performance"). In those cases, "the act fails to be *safely* successful, since it might too easily have failed, through lack of the required competence or conditions. It might still be apt, nevertheless, indeed attributably, creditably apt" (Sosa, *A Virtue* 81). So, Sosa's account give us the right results when considering the *adapted* fake-barn cases because it ignores safety, but the wrong ones when considering the *original* case because its aptness condition is satisfied (cf. Greco, *Achieving*). After all, in the original case, the subject isn't using the right "tool" for those conditions, although she happens to get the truth. Nevertheless, the fact that she could have very easily used some other tool, without being epistemically irresponsible, when she actually used the right one doesn't undermine the actual success.

Grabit case (Harman 142-4), which are sometimes (misleadingly) thought to exploit the "social aspect" of knowledge (Meeker). In these cases, it seems, according to Gilbert Harman, that the subject doesn't have knowledge and this seems to be the result of there being evidence that can easily be accessed by the subject, which isn't anyway acting irresponsibly (i.e. being negligent with regard to the evidence) and which, if so accessed, it would count as a psychological defeater. For example, in the assassination case, one has gained a true belief about the death of the dictator by testimony (say, some newspaper), but when one is isolated from media reports in one's hotel room, false denials of the dictator's death are being aired. But I am not sure most people would share Harman's intuitions (I, for one, don't) and indeed some claim that many don't (Lycan 121).

Moreover, according to our account, the subject does know in these cases. This is because the knowledge-yielding procedures are silent about this possibility of luckily, but not irresponsibly, lacking some misleading evidence that is easily accessible and that would undermine knowledge, so no condition is being violated in these cases. And this seems to be as it should, given the resemblance of this case to cases of benign evidential luck–when it is lucky that the agent acquires the evidence that she has in favour of her belief. Although in these cases the luck doesn't concern the evidence that becomes available but the evidence that doesn't, still we are lucky, but not irresponsible, with regard to the total evidence acquired. And this kind of luck doesn't seem knowledge-precluding.

In any case, given the controversial nature of the intuitions behind Harman cases, they seem to lack the counter-exemplary force that Gettier cases and Goldman-Ginet cases enjoy (Lycan 125). Because of this, Harman cases are unlikely to trump theories. Nevertheless, if we are right about these cases instantiating a sort of benign evidential luck, we can understand why many would be wary of the denial of knowledge in these cases.

4.3 Local and Global Reliabilism

Finally, note that the account proposed, as a form of *global reliabilism*, can eliminate cases of what we might call *procedural veritic luck*. This is the kind of knowledge-undermining luck that sensitivity and safety theories suffer from, because of their shared *local reliabilism*, where the method employed need only be reliable with regard to the targeted proposition (Goldman, *Epistemology*; McGinn). Let me explain.

A local reliabilist theory of knowledge, as Colin McGinn says, "localizes the conditions for knowledge into a relation between the knower and a unique proposition [whereas a global reliabilist theory] speaks of the person's propensity to believe the truth with respect to a range of distinct 'relevant' propositions" (116-7). So, the main difference between a local and a global reliability theory concerns the range of propositions taken into account or, as Goldman says, "uses for which the process is reliable"

(*Epistemology* 44-5). Indeed, "Global reliability is reliability for all (or many) uses of the process, not just its use in forming the belief in question" (*Epistemology* 45). Roughly put, the difference between these theories lies in the emphasis on the reliability of either type processes (global reliabilism) or particular propositions (local reliabilism). More precisely, we need to distinguish between the reliability of a process in general and the reliability of a process with respect to a given proposition.

The relevance of this distinction is that a local account, such as the sensitivity or safety theories, is then susceptible to counterexamples that exploit their locality. But before making this clear, let me briefly introduce these local theories. A sensitivity theory states that knowledge is sensitive true belief, where this sensitivity is understood as the satisfaction of the following subjunctive: not-p→not-B(p); which in a crude but intuitive reading states that "if *p* were false, S wouldn't believe that *p*" and in evaluating this condition we consider only close possible worlds (Nozick).[15] A safety theory, on the other hand, states that following subjunctive: B(p)→p; which again in a crude, but intuitive reading, states that "if S believed that *p*, then *p* would be true" and, again, only close possible worlds are relevant to its evaluation (Sosa, "How to Defeat").

Now, needless to say, both principles have been criticized for a variety of reasons.[16] But regardless of the success of these criticisms, safety theories can't rule out cases involving procedural luck. That is, even if these theories might naturally be seen as arising out of the need to be able to eliminate cases of malign luck (Pritchard, *Epistemic Luck*, "Knowledge, Luck"), because of their local nature, they can't rule out the sort of luck that global theories can. After all, globally unreliable procedures that safely or sensitively deliver the truth, say by some quirk of nature, for the targeted propositions seem possible. Intuitively, it seems that the subject doesn't know in these cases, precisely because the procedure employed isn't globally reliable – it is in a sense lucky that the belief is true and safe, since the procedure isn't globally reliable. The local reliability of a procedure shouldn't be happenstance and the above sort of scenario hasn't gone unnoticed. For example, Pritchard (*Knowledge*, "Knowledge and Understanding") introduces a case where one can safely acquire the truth because some person makes sure that one does so with regard to the targeted proposition (by arranging the world accordingly) in most nearby (and in all very close nearby) possible worlds (hence locally reliable), but the procedure employed isn't truth-conducive across a range of relevant propositions.[17]

---

[15] Strictly speaking Nozick holds a stronger theory which requires a fourth ("adherence") condition: p→B(p). So his "tracking" theory requires subjunctive sensitivity to both truth and falsehood; while "sensitivity theories" require only subjunctive sensitivity to falsehood.

[16] See e.g. Comesaña; Greco "Worries;" Hiller and Neta; Pritchard, *Epistemic Luck*; Sosa, "How to Defeat," "Skepticism."

[17] But, couldn't we use this helper-scenario to create counter-examples to global accounts, including the one here proposed? Notice that, although a range of our beliefs would in this case be true (the helper

So in response to this deficiency of local reliabilism one might attempt to combine local and global reliabilism, as anti-luck virtue epistemologists do (Pritchard, "Knowledge and Understanding"). In this case, one would give up the pure safety or sensitivity account for a hybrid account that exploits both local and global features—say, safety and virtue-reliabilist conditions, respectively. There is indeed no problem in mixing these two approaches since they aren't mutually exclusive (Goldman, *Epistemology* 47).[18] And in fact this sort of hybrid approach is already attempted by Goldman (*Epistemology*), who endorses both local and global reliability (cf. McGinn). But if what I say above is correct, no such hybrid is required. Our version of global reliabilism can preclude the sort of veritic double-luck involved in Gettier cases and the environmental luck involved in Goldman-Ginet cases, as well as the above procedural luck. No supplementation from local principles, such as safety and sensitivity, is required. Having said that, the pure local alternative to understanding reliability is unsatisfactory, if it is not supplemented by a global condition. That is, pure safety and sensitivity theories, where the subject's belief is "subjunctively connected to the fact" (Nozick 178), will not do, since they fail to rule out cases of procedural luck.

To sum up, there are different kinds of epistemic luck and not all are knowledge-undermining. But veritic double-luck, environmental-luck and procedural-luck are epistemically deleterious. Now pure local reliabilist accounts can't do justice to procedural luck, since a procedure can be locally reliable by accident, hence suggesting the adoption of hybrid theories that combine local and global reliabilism. But this isn't required since the global reliabilist account here proposed can do justice to the different sorts of knowledge-threatening veritic luck; hence satisfying, I suggest, the non-accidentality desideratum.

Of course I anyway anticipate costs to this account. Even if our account can handle Gettier cases and the like, and so capture the non-accidentality desideratum, it might fail to capture some other desideratum, in which case we would have to weight the costs and benefits of the competing theories to adjudicate between them. But regardless of that, this account seems to possess the means to capture a desideratum that others don't.

5 Conclusion

Fifty years on since Gettier's article, there is I think one lesson to be learned: one can make progress in epistemology if one does not ignore the social dimensions of our prac-

makes sure they are), the procedure itself doesn't connect us to the facts that make the targeted propositions true (that is indeed what the helper is doing for us); hence depriving us of knowledge.

[18] Although one might worry the move is *ad hoc*. Pritchard ("Knowledge and Understanding" 59-62) attempts to address this.

tices. Indeed, as Jonathan Kvanvig says, "we should never begin to think that the deepest epistemological questions concern the isolated intellect" (*The Intellectual Virtues* 177). The account here proposed, by positing a reflective endorsement of the knowledge-yielding procedures that can be met at the social level, grants a fuller engagement with the sort of social issues to which traditional epistemology is often blind given its strongly individualist orientation. This social approach corrects the unfortunate individualistic simplifications of much current mainstream epistemology which are inapt for theorizing about knowers who are members of social communities and which won't allow us to capture the perspectival dimension on knowledge. This correction, I have argued, allows us to capture the non-accidentality desideratum, showing that the pessimism displayed by some against the project of adding some accidentality-blocker to true belief to have knowledge is misplaced.

References

1. Ayer, Alfred. "Comments on Professor Williams' "Knowledge and Reasons"." *Problems in the Theory of Knowledge*. Ed. Georg Henrik von Wright. The Hague: Nijhoff, 1972. 20-35. Print.
2. BonJour, Laurance. *The Structure of Empirical Knowledge*. MA: Harvard University Press, 1985. Print.
3. Chisholm, Roderick. *Theory of Knowledge*. NJ: Prentice Hall, 1977. Print.
4. Comesaña, Juan. "Unsafe Knowledge." *Synthese* 146 (2005): 395-404. Print.
5. Dancy, Jonathan. *Introduction to Contemporary Epistemology*. Oxford: Blackwell, 1985. Print.
6. De Brasi, Leandro. "Reliability and Social Knowledge-Relevant Responsibility." 2013. TS.
7. De Brasi, Leandro. "Testimony and Value in the Theory of Knowledge." 2013. TS.
8. Fogelin, Robert. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press, 1994. Print.
9. Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 23 (1963):121-3. Print.
10. Goldman, Alvin. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73 (1976): 771-91. Print.
11. Goldman, Alvin. *Epistemology and Cognition*. MA: Harvard University Press, 1986. Print.
12. Greco, John. *Putting the Skeptics in Their Place*. Cambridge: Cambridge University Press, 2000. Print.
13. Greco, John. "Worries about Pritchard's Safety." *Synthese* 158 (2007): 299-302. Print.
14. Greco, John. *Achieving Knowledge*. Cambridge: Cambridge University Press, 2010. Print.
15. Grimm, Stephen. "Understanding." *The Routledge Companion to Epistemology*. Eds. Sven Bernecker, Duncan Pritchard. London: Routledge, 2011. 84-94. Print.
16. Harman, Gilbert. *Thought*. NJ: Princeton University Press, 1973. Print.
17. Hiller, Avram and Neta, Ram. "Safety and Epistemic Luck" *Synthese* 158 (2007): 303-13. Print.

18. Howard-Snyder, Daniel, Howard-Snyder, Frances, and Feit, Neil. "Infallibilism and Gettier's Legacy." *Philosophy and Phenomenological Research* 66 (2003): 304-27. Print.
19. Hyman, John. "How Knowledge Works." *Philosophical Quarterly* 49 (1999): 433-51. Print.
20. Kvanvig, Jonathan. *The Intellectual Virtues and the Life of the Mind*. Maryland: Rowman & Littlefield, 1992. Print.
21. Kvanvig, Jonathan. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press, 2003.
22. Lehrer, Keith. *Theory of Knowledge*. Boulder: Westview Press, 1990. Print.
23. Lewis, David. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (1996): 549-67. Print.
24. Lycan, William. "Evidence One Does Not Possess." *Australasian Journal of Philosophy* 55 (1977): 114-26. Print.
25. McGinn, Colin. *Knowledge and Reality*. Oxford: Clarendon, 1999. Print.
26. Meeker, Kevin. "Justification and the Social Nature of Knowledge." *Philosophy and Phenomenological Research* 69 (2004): 156-72. Print.
27. Millar, Alan. "Knowledge and Recognition." *The Nature and Value of Knowledge*. Duncan Pritchard, Alan Millar, Adrian Haddock. Oxford: Oxford University Press, 2010. 91-190. Print.
28. Nozick, Robert. *Philosophical Explanations*. MA: Harvard University Press, 1981. Print.
29. Pritchard, Duncan. *Epistemic Luck*. Oxford: Oxford University Press, 2005. Print.
30. Pritchard, Duncan. "Knowledge, Luck and Lotteries." *New Waves in Epistemology*. Eds. Vincent Hendricks, Duncan Pritchard. Hampshire: Palgrave Mcmillan, 2008. 28-51. Print.
31. Pritchard, Duncan. *Knowledge*. Hampshire: Palgrave, 2009a. Print.
32. Pritchard, Duncan. "Apt Performance and Epistemic Value." *Philosophical Studies* 143 (2009b): 407-16. Print.
33. Pritchard, Duncan. "Knowledge and Understanding." *The Nature and Value of Knowledge*. Duncan Pritchard, Alan Millar, Adrian Haddock. Oxford: Oxford University Press, 2010. 3-90. Print.
34. Roush, Sherrilyn. *Tracking Truth*. Oxford: Clarendon, 2005. Print.
35. Russell, Bruce. "Epistemic and Moral Duty." *Knowledge, Truth and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press, 2001. 34-48. Print.
36. Sosa, Ernest. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13 (1999): 141-53. Print.
37. Sosa, Ernest. "Skepticism and Contextualism." *Philosophical Issues* 10 (2000): 1-18. Print.
38. Sosa, Ernest. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Vol. 1. Oxford: Clarendon, 2007. Print.
39. Sturgeon, Scott. "The Gettier Problem." *Analysis* 53 (1993): 156-64. Print.
40. Unger, Peter. "An Analysis of Factual Knowledge." *Journal of Philosophy* 65 (1968): 157-70. Print.
41. Williams, Bernard. "Knowledge and Reasons. " *Problems in the Theory of Knowledge*. Ed. Georg Henrik von Wright. The Hague: Nijhoff, 1972. 2-19. Print.
42. Williams, Michael. "Responsibility and Reliability." *Philosophical Papers* 37 (2008): 1-26. Print.

43. Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000. Print.
44. Zagzebski, Linda. "The Inescapability of Gettier Problems." *Philosophical Quarterly* 44 (1994): 65-73. Print.
45. Zagzebski, Linda. *Virtues of the Mind*. Cambridge: Cambridge University Press, 1996. Print.
46. Zagzebski, Linda. "What is Knowledge?" *The Blackwell Guide to Epistemology*. Ed. John Greco. Oxford: Blackwell, 1999. 92-116. Print.
47. Zagzebski, Linda. *On Epistemology*. Belmont: Wadsworth, 2009. Print.