NO LUCK IN THE DISTANCE: A REPLY TO FREITAG

Fernando Broncano-Berrocal

(published in *Theoria: A Swedish Journal of Philosophy*)

Abstract

In a recent article in this journal, Wolfgang Freitag argues that Gettier-style cases that are based on the notion of "distant" epistemic luck cannot be ruled out as cases of knowledge by modal conditions such as safety or sensitivity. I argue that (1) safety and sensitivity can be easily fixed and that (2) Freitag provides no convincing reason for the existence of "distant" epistemic luck.

In a recent article in this journal, Wolfgang Freitag argues that Gettier-style cases that are based on the notion of *distant* epistemic luck, i.e., "luck which discloses itself in distant possible worlds alone" (Freitag 2014, p. 45), cannot be ruled out as cases of knowledge by modal conditions such as

> SAFETY (if S knows that p, then S believes that p truly in the actual world, and in nearly all, if not all, close possible worlds in which the conditions for belief formation are the same as in the actual world, S only believes that p when p is true)[1]

[1] Adapted from Pritchard (2005, p. 163). See Broncano-Berrocal (2014) for a defense of the safety principle against counterexamples to its necessity for knowledge.

or

> SENSITIVITY (if S knows that p, then in the closest possible worlds where p is false and in which the conditions for belief formation are the same as in the actual world, S does not believe that p).[2]

In this essay, I will argue that:

1. Safety and sensitivity can be easily fixed just by adding a clause to the effect that one must not believe closely related false propositions.

2. Freitag provides no convincing reason for the existence of "distant" epistemic luck. The categorization of some examples as cases of distant luck has to do with a confusion between the degree of inexorability of the conditions that enable belief formation and the degree of inexorability of the conditions that bring about the fact that one gets things right or wrong.

The relevance of this essay is twofold: it not only offers a defense of the safety and sensitivity conditions against Freitag's objections, but also of the widespread assumption in epistemology that knowledge-undermining possibilities are close possibilities. If distant epistemic luck existed, that assumption would be wrong, and as a consequence accounts of epistemic luck and of knowledge based on that assumption, as well as the solutions to epistemological problems offered by the latter, would have to be critically reconsidered. This essay will give arguments to avoid such a revisionary program.

[2] Adapted from Nozick (1981, p. 179).

1 Safety and sensitivity can be easily fixed

The problem for the safety and sensitivity conditions for knowledge is the following. As Freitag acknowledges, safety and sensitivity get support from the fact that they are able to rule out Gettier-style cases as cases of knowledge. To the extent that Gettier-style cases are paradigmatic cases of knowledge-undermining luck, it is fair to say that safety and sensitivity exclude dangerous epistemic luck. That constitutes excellent motivation for using safety or sensitivity in one's account of knowledge. But if there were a case in which there is no knowledge, because of being a Gettier-style case (i.e., a case of luck), but in which safety and sensitivity nevertheless hold, the case would show that they are not anti-luck conditions after all, or at least that they are not adequate anti-luck conditions, which would in turn undermine the original motivation for including them in one's account of knowledge. Freitag thinks to have found such a case:

> CLOCK
>
> [A] clock malfunctions (...) and shows either 3:00 or 6:00. It shows 3:00 at 3:00 and at all times between 6:00 and 11:58 (a.m. and p.m.). At all other times, it shows 6:00. Jim, not aware of the clock's malfunctioning, looks at the clock at 3:00, thereby picking up the true belief that it is 3:00 (Freitag 2014, p. 54).

CLOCK seems to be a Gettier-style case in the same way as Russell's original clock case is (Russell 1948, pp. 170-171). In the latter case, Jim does not know that it is 3:00 because he happens to look at 3:00 at a clock that he justifiably believes to be very reliable but that stopped 24 hours before. One difference is that in Russell's case the clock is stopped, while in CLOCK it malfunctions. For discussion purposes, however, that difference is irrelevant, insofar as both cases involve unreliable ways to form beliefs about the time.

According to Freitag, the crucial difference between the cases lies in the kind of luck instantiated. While Russell's case entails that Jim would form the false belief that it is 3:00 *in close possible worlds*, in CLOCK he would believe that falsely only in *distant* (i.e., *non-close*) *possible worlds*. As usually conceived, closeness between possible worlds is fixed by intuitive similarity. Freitag's idea is that the relevant metric of similarity is given by context. In CLOCK and in Russell's case, the metric is determined by the time at which Jim forms his belief. In this way, he thinks that worlds in which Jim looks at the clock at 2:59 or at 3:01 are close. By the same token, possible worlds in which Jim forms the false belief that it is 3:00 when looking at the clock at 6:00 count as distant, according to Freitag.

The above formulations of safety and sensitivity can be easily fixed once we notice that, although in close possible worlds Jim would not form a false belief in the same proposition that he actually believes (namely, that it is 3:00), he would believe a closely related false proposition about the time (namely, that it is 6:00). In other words, safety and sensitivity can be easily fixed once we notice that in CLOCK the relevant knowledge-undermining possibility is not the close possibility of believing p falsely when one actually believes p truly, but the *close possibility of believing closely related false propositions*.[3] For the sake of clarity, let me give some examples in which one's actual knowledge of a proposition is undermined because one could have easily believed closely related propositions falsely.

The first example is suggested by Plato in the *Theaetetus* (207e-208a). Theaetetus agrees with Socrates that, when at such a stage in her progress, a person in writing "Theaetetus" actually writes the name with "Th" and "e", but writes or could have easily written "Theodorus" with "T" and "e", then that person does not know the first syllable of "Theaetetus". In the same way, if under normal light conditions one comes to believe that a

[3]See Manley (2007) for relevant discussion on the different kinds of knowledge-undermining possibilities, including the close possibility of believing closely related false propositions.

wall painted red is red, but also believes or could have easily believed that the opposite wall painted with a very similar shade of red is not red, then one does not seem to know that the former wall is red. On the other hand, if one reasons that 3 is prime because it has no other positive divisors than 1 and itself, but one also believes or could have easily believed that 2 is not prime after having reasoned in the same way, then one hardly knows that 3 is prime. Finally, if one believes that it is 3:00 when looking at a malfunctioning clock at 3:00, but could have easily believed a different false proposition about the time using that clock, such as the proposition that it is 6:00 when looking at the clock at 3:01, then one does not know that it is 3:00. The same applies to propositions concerning how many miles a car has run when consulting a malfunctioning odometer (see Freitag's odometer case below).

The fact that certain propositions are "closely related" has to do with the way our belief-forming capacities work. Note the following: one neither forms the belief that number three is prime by using one's taste, nor does one form the belief that the ice cream one is eating is tasty by competent deduction. The general thought is that most of our cognitive abilities and methods of belief formation are domain-specific, which means in turn that they tend to be relativized to fields of propositions, e.g., propositions concerning shape, distance, color and so on in the case of visual beliefs (it is a task of cognitive science to specify the criteria for individuating the relevant fields of propositions in each case). It seems that if several propositions are in the field to which a cognitive ability or belief-forming method is relativized, those propositions automatically count as closely related, to the extent that if one comes to know a proposition in that field by that cognitive ability or belief-forming method, at a minimum the ability or method must not easily lead (or have led) one to believe propositions in that field falsely, on pain of losing one's knowledge. Although that might not generalize, it is at least true of some relevant cases, including

Freitag's.[4]

Safety and sensitivity can be easily fixed just by adding a clause to the effect that one must not form false beliefs in closely related propositions in the field of propositions to which one's belief-forming method is relativized:

> SAFETY: If S knows that p, then S believes that p truly in the actual world, and in nearly all, if not all, close possible worlds in which the conditions for belief formation are the same as in the actual world, S only believes that p when p is true, and S does not believe false propositions closely related to p.[5]

[4]I do not claim that *all* our belief-forming capacities are relativized to fields of propositions. For example, one can know a single proposition by testimony even though one also has many false beliefs in related propositions. What I claim is that most of our belief-forming capacities have that feature, and that is sufficient to give rise to the special kind of knowledge-undermining possibility distinguished here. In any case, and to reply to a possible ad-hocness objection, the idea that our cognitive abilities and belief-forming methods are relativized to fields of propositions is anything but new: it is accepted by anyone upholding a global reliability condition on knowledge, that is, an epistemic condition to the effect that if one knows that p, one's belief that p is the output of a belief-forming process or method that is reliable not only with respect to p but to a wider range of propositions. Most reliabilists (including proper functionalists and virtue reliabilists) accept this rather uncontroversial requirement (uncontroversial among externalists, of course). See Goldman (1986, Ch. 3), Graham (*forthcoming*), Greco (2000, 216), and Sosa (1991, Ch. 16) for relevant discussion.

[5]In his (2005), Pritchard understands safety as the requirement that one knows that p only if one believes that p truly in the actual world and one does not form false beliefs *in* p on the same basis in close possible worlds. More recently, he has changed his view: "[N]otice that this is not the version of safety which is generated by anti-luck epistemology. Our interest is instead more broadly on whether the same actual basis for belief leads to false belief in close possible worlds. Any false belief formed in close possible worlds on the same basis is thus relevant to the safety of the target belief, and not just false beliefs formed in the very same proposition as in the actual world" (Pritchard 2015: 102). However, if *any* false belief formed in close possible worlds on the same basis can make safety fail, safety is not necessary for knowledge. Suppose that by mathematical reasoning one comes to believe that 2 is a prime number in the actual world and that, on the same basis, one forms the false belief that 4,035,540,087 is a prime number in close possible worlds. It is not unnatural to claim that one actually knows that 2 is prime despite the easy possibility of believing the latter proposition, but on Pritchard's new view, one's belief counts as unsafe. Cases of this sort motivate the introduction of the clause on closely related propo-

SENSITIVITY: If S knows that p, then in the closest possible worlds where p is false, and in which the conditions for belief formation are the same as in the actual world, S neither believes that p nor false propositions closely related to p.

Consider CLOCK again. In close possible worlds, Jim believes a false proposition about the time ("It is 6:00") that is closely related to the true proposition he actually believes ("It is 3:00"). The same occurs in the closest possible worlds in which it is not 3:00. That is why Jim's belief is, respectively, unsafe and insensitive.

2 Why is it important to make error possibilities explicit?

When it comes to giving counterexamples to modal principles like safety or sensitivity, it is very important that the way the cases are described leaves no room for ambiguous interpretations of what the relevant error possibilities are. In order to see this, consider the other Gettier-style counterexample given by Freitag, which is allegedly based on distant epistemic luck as well:

ODOMETER

Jenny, who knows next to nothing about cars, is looking for a second-hand car. At the dealer's she has a closer look at an older Pontiac whose odometer shows 68,500 miles. Jenny therefore (and for no other reason) believes the car to have run less than 69,000 miles. In the evening she discusses the alleged bargain with her best friend. In response to her friend's doubts,

sitions. Presumably, while the false proposition that 4 is a prime number is closely related to the proposition that 2 is a prime number, the false proposition that 4,035,540,087 is a prime number is not. In this way, believing the latter proposition in close possible worlds does not undermine the safety of one's actual belief that 2 is a prime number, but that is not the case if the relevant error possibility concerns the proposition that 4 is prime.

> Jenny insists that she knows the car to have run less than 69,000 miles. Indeed, the former owner barely drove the car and so the car really has run only 68,500 miles, the odometer showed the correct number, and Jenny's belief is true. But now suppose that she has looked at a defective odometer, whose last four digits work perfectly, but whose first two digits are stuck on "06". The odometer presents any number of miles up to 69,999 miles correctly; with the 70,000th mile, however, the odometer rolls over to 60,000. Ignorant of this defect, Jenny does not take this possibility into account (Freitag 2014, pp. 57-58).

Freitag's intuition is that Jenny does not know that the car has less than 69,000 miles. I will argue that if that intuition is correct, it is because Jenny's belief is unsafe and insensitive after all. The reason why Freitag thinks that safety and sensitivity cannot rule out ODOMETER as a case of knowledge is the following. On the one hand, Jenny's belief that the car has run less than 69,000 miles is safe because there are no close possible worlds in which Jenny believes that proposition falsely (according to Freitag (2014, p. 58), the relevant world ordering is fixed by how many miles the car has run: "In 68,400-worlds, as in 68,600-worlds, Jenny continues to hold her belief and this belief is still correct. The worlds in which Jenny believes falsely, via the same method, that the car has run less than 69,000 miles, e.g., 70,000-worlds, do not, by my reckoning, count as relevantly close"). On the other hand, Freitag thinks that Jenny's belief is sensitive because in the closest possible worlds in which the car has run more than 69,000 miles (by Freitag's standards, these are 69,001-worlds), Jenny does not believe the car to have run less than 69,000 miles (recall that up to 69,999 miles the odometer works well).

The description of ODOMETER does not specify how easily the former owner could have driven the car more than 70,000 miles, and hence how easily the odometer could have malfunctioned thus producing a false be-

lief in Jenny.[6] In what follows, I will try to show that the ignorance intuition arises if and only if one implicitly fills in the missing details with some story to the effect that such things could have been easily the case, but then there is no reason to keep thinking that Jenny's belief is safe and sensitive. To that aim, I will give two interpretations of ODOMETER: one in which there is a clear ignorance intuition, but in which the relevant error possibilities are close; another in which it is clear that the relevant error possibilities are distant, but in which the ignorance intuition vanishes.

The following is one a way of specifying ODOMETER among the many structurally equivalent specifications that would give rise to a clear ignorance intuition (ODOMETER 1). Suppose that the former owner wanted to become a member of a Pontiac club whose regulations say that only owners of Pontiac cars with more than 70,000 miles can be members of the club. Suppose that while she was embarked on a road trip trying to reach the 70,000 miles goal, her company's stock dropped 99% due to an unexpected financial scandal leading her to bankruptcy and making her sell her Pontiac at the first car dealer along the road.

Do 70,000-worlds count as distant under that specification of the case? If we make vivid the fact that the scandal could easily have broken out at any time of the day (or the week) in the description of the case (so that it could have easily turned out that the former owner would have reached the 70,000 miles goal), then 70,000-worlds will be surely part of the set of close possible worlds. But if that is so, then Jenny's belief is not safe, because she could have easily believed that the car has run less than 69,000

[6] In the same way, Freitag's description of CLOCK leaves unspecified the time frame in which Jim could easily have consulted the clock. All we know in this regard is the following: "I assume the clock's hands to be the only evidence the epistemic agent can rely on in her estimation of the time; assume the respective agents to be on a winter expedition close to the North Pole" (Freitag 2014, p. 51). Thus, it is left unspecified whether Jim could have easily looked at the clock at a time close to 3:00 (e.g., at 2:59 or at 3:01) or within the whole span of time of the polar night, which lasts for more than 24 hours.

miles falsely. Jenny's belief is not sensitive either, because according to the new contextually given standards 70,000-worlds are among the closest possible worlds in which the car has run more than 69,000 miles, and those are possible worlds in which she would continue to believe that the car has run less than 69,000 miles.

Now consider an interpretation of ODOMETER (ODOMETER 2), which is consistent with Freitag's description of the case, and in which it is practically impossible for the former owner to reach the 70,000 miles goal, so that the possibility of forming the false belief that the car has run less than 69,000 miles is anything but close. Suppose for instance that the odometer is connected to a bomb that would explode if the car runs for more than 69,999 miles. If we take that into account, it is not so clear that ODOMETER is not a case of knowledge. For while in CLOCK the clock is not a reliable indicator of the time when Jim forms his belief, in ODOMETER the odometer reliably indicates the number of miles when Jenny forms hers, and it is not the case that it could easily stop being a reliable indicator of the traveled distance. When the case is so described, Jenny's belief counts as safe and sensitive.[7]

[7] ODOMETER 1 is structurally *equivalent* to a modified version of a case discussed by Alan Goldman (1987: 183-4) and due to Risto Hilpinen (call it THERMOMETER 1): Robert looks at a thermometer that is accurate within the range of 0 to 100 degrees F. At all temperatures below 0 the thermometer registers 0 degrees. By observing its reading of 0 degrees (it is actually 0 degrees), Robert comes to believe that it is 0 degrees outside. The temperature might very easily have oscillated till 5 degrees down. Robert's belief is unsafe (because too easily could his belief have been false) and insensitive (because he would still believe that it is 0 degrees outside even if it were not so). On the other hand, ODOMETER 2 is structurally very *similar* to a version of THERMOMETER 1 (call it THERMOMETER 2), in which (1) the actual temperature is 70 degrees F, (2) the possible oscillation of temperature is the same, and (3) Robert comes to believe that it is not -50 degrees outside. Goldman, for instance, has the intuition that THERMOMETER 2 is a case of knowledge (of the proposition that it is not -50 degrees outside). He also presents the case as a counterexample to the necessity of sensitivity, because if it were the case that it is -50 degrees outside, Robert would not believe it (the thermometer would read "0 °F") (ODOMETER 2 is similar but not equivalent to THERMOMETER 2 in that Jenny's actual belief is sensitive because possible worlds in which the car runs for more than 69,999 miles are worlds in which she would not form a false belief, because the bomb would explode,

What these two different possible interpretations of ODOMETER show is not only that Freitag's unspecific description makes his case a contentious counterexample to safety and sensitivity, but more importantly, that once we specify the case further, thus shedding light on the possible structures it might have, the problem dissolves. If we fill in the details of the case such that the ignorance intuition clearly arises, the relevant error possibilities count as close and Jenny's actual belief counts as unsafe and insensitive. On the other hand, if we fill in the details such that the relevant error possibilities are clearly distant and the belief is clearly safe and sensitive, the ignorance intuition vanishes. Either way, ODOMETER is no counterexample to safety and sensitivity. In sum, the unspecific description of the case is what makes Freitag think that the ignorance intuition can go hand in hand with the satisfaction of safety and sensitivity. Further specifications of the case, however, prove Freitag wrong.

3 No reason to think that "distant" epistemic luck exists

Finally, Freitag provides no convincing reason to think that his cases instantiate "distant" epistemic luck. In order to understand this point, it is important to take a step back and think about what makes us deem an error possibility (e.g., the possibility that one believes p falsely when one actually believes p truly) knowledge-undermining. A standard requirement in the epistemological literature is that a possibility must be close

the former owner would die and the car would not have been around for Jenny to inspect). Counterexamples to the necessity for knowledge of sensitivity are well-known, but as far as the purposes of the safety-theorist are concerned, there is an important insight to be gained by comparing the two versions of ODOMETER with THERMOMETER 1 and 2: once the relevant cases are placed in a continuum, safety-theorists can reject Freitag's diagnosis of ODOMETER and similar cases, to wit, the closer the cases are to cases of type 1, the stronger the intuition that agents do not know, but at the same time the less safe their beliefs are; on the other hand, the closer the cases are to cases of type 2, the stronger the intuition that agents know, but at the same time, the safer their beliefs are. Either way, knowledge and safety do not come apart. Thanks to an anonymous reviewer for suggesting this point.

(i.e., sufficiently similar in the relevant respects) in order for it to count as knowledge-undermining. Although that is of course what is in dispute (that is the challenge posited by Freitag to a widespread assumption in epistemology), the requirement is meant to explain at least why some bizarre possibilities do not undermine ordinary knowledge. For instance, one knows that one is drinking water (H_2O) even though in some possible world one is drinking XYZ and mistakenly thinks it is water. The reason why such a possibility does not undermine one's knowledge is that the Twin Earth possibility is *very* far away from actuality.[8]

There is no reason to think that such an uncontroversial requirement is incompatible with the existence of distant epistemic luck. After all, the error possibilities that exemplify the phenomenon of distant luck, although (allegedly) non-close in a sense, are not so far away (i.e., not so dissimilar) as bizarre possibilities. In other words, if there is such a thing as "distant" epistemic luck, the relevant luck does not disclose itself in so distant possible worlds. Still, its existence somehow undermines the position of any epistemologist who grounds his or her theory on the closeness assumption. In what follows, it is not my intention to argue for the consequent of the former conditional. I will rather argue that Freitag provides no convincing reason for the truth of the antecedent.

When assessing whether a possibility is knowledge-undermining or not, one should always try to draw as clearly as possible the distinction between the degree of inexorability of the *conditions that enable the formation of the belief* (i.e., how easily those conditions could have been different) and the degree of inexorability of the *conditions that bring about the fact that one gets things right or wrong* —call them the "determining conditions"— (i.e., how easily one could have formed a belief with different content or truth value in the circumstances). Both things are often conflated, but only the

[8] A different issue is that of not taking radical hypotheses seriously because of being far away. Although that might be a solution to Cartesian skepticism, it is a cheap one. As such, alternative considerations might be needed.

latter seems relevant to judge the epistemic status of a belief.

By way of illustration, suppose that one is in a room with only one door. Suppose that one opens the door, takes a look at the well-functioning clock behind it and as a result forms the true belief that it is 3:00. The opening of the door merely enables one to form the belief, while the fact that there is a well-functioning clock reading "3:00" together with the fact that one places one's trust on it is what brings about that one gets things right. Plausibly, one's belief amounts to knowledge.

Now consider the following specification of CLOCK, which is consistent with Freitag's description of the case (call it CLOCK 1). In CLOCK 1, Jim faints, loses all sense of time, wakes up and takes a look at the clock. Suppose that the description of the case makes it vivid that, given Jim's condition, he could have easily woken up at any time (for all we know given how Freitag introduces CLOCK, Jim could have been unconscious for several days). Suppose that luck has it that he wakes up at 2:59 and takes a look at the clock at 3:00 thereby forming the true belief that it is 3:00. The relevant error possibilities come out close and, as a result, the fact that Jim does not know that is 3:00 is not due to distant luck. The following case is structurally equivalent to CLOCK 1:

> DOORS
>
> Jim is in a room with sixty doors and does not know what time it is. Unbeknownst to him, each door hides a malfunctioning (an unreliable) clock. Half of them read "6:00"; the other half read "3:00". It is 3:00. Jim randomly picks one of the doors, opens it, takes a look at the clock before him and as a consequence forms the true belief that it is 3:00.

I surmise that most people will have the intuition that Jim does not know that it is 3:00. This is explained by the fact that the very opening of the door, which in an ordinary clock situation would figure as part of the con-

ditions enabling belief formation, is part of the conditions that help bring about that Jim gets things right, which is why the relevant error possibilities come out close and, therefore, why distant luck does not explain Jim's ignorance of the fact that it is 3:00.

As pointed out before, only determining conditions are relevant to assess the epistemic status of a belief. But Freitag's description of CLOCK, by leaving unspecified the time frame in which Jim could *easily* have consulted the clock, makes impossible to judge where the enabling conditions end and where the determining conditions start (similar considerations apply to ODOMETER).

Suppose that we go in the opposite direction and we specify CLOCK so that Jim could have consulted the malfunctioning clock only at times close to 3:00 (e.g., at 2:59 or at 3:01). Call this version CLOCK 2. A structurally equivalent case would be the following:

> LOCKED DOORS
>
> Jim is in a room with sixty security doors and does not know what time it is. Unbeknownst to him, each single door hides a malfunctioning (an unreliable) clock. Half of them read "6:00"; the other half read "3:00". All doors hiding clocks reading "6:00" are open. All doors hiding clocks reading "3:00" are locked in a way that it is physically impossible for Jim to open them, except for one, which is unlocked exactly at 3:00. It is 3:00. Jim, decided to get out of the room, randomly picks one of the doors, opens it, takes a look at the clock before him and as a consequence forms the true belief that it is 3:00. Luckily, he has opened the only door with a clock reading "3:00" that was not locked.

LOCKED DOORS is structurally equivalent to CLOCK 2 for the following reasons: in both cases (1) Jim truly believes that it is 3:00 in the actual

world, (2) close possible worlds are worlds in which Jim believes that it is 6:00 falsely (in LOCKED DOORS in particular, this is due to the fact that Jim could easily have opened any of the doors hiding clocks reading "6:00" at 3:00 or at times close to 3:00), and (3) distant worlds are worlds in which Jim believes that it is 3:00 falsely (in LOCKED DOORS, those are worlds in which Jim somehow manages to open one of the locked doors). That being so, if there is no knowledge in CLOCK 2 because of knowledge-undermining *distant* epistemic luck, there should also be no knowledge in LOCKED DOORS for the same reason.

However, that is an unwelcome consequence for Freitag. According to Freitag, distant epistemic luck is luck that discloses itself in distant possible worlds *alone*. Therefore, his reason to think Jim's actual belief is luckily true in LOCKED DOORS and hence his explanation of why Jim lacks knowledge would be that there are distant worlds in which Jim manages somehow to open one of the locked security doors consequently forming a false belief in the same proposition he believes in the actual world. But why do we have to accept that the epistemic luck in play must be explained in terms of what happens in such distant worlds, provided that in virtue of how things stand in the actual world it is physically impossible for Jim to open the locked doors? That explanation is, if not wrong, at least strained.

Note that Freitag cannot use DOORS or the corresponding specification of CLOCK (i.e., CLOCK 1) to make the point that distant epistemic luck exists. As I have argued, in both cases the relevant error possibilities count as close, so the fact that the target beliefs do not amount to knowledge cannot be explained in terms of distant luck (but one can explain the absence of knowledge by pointing to the unsafety or to the insensitivity of the beliefs). The failure of that possible move by Freitag together with the wrong result in the explanation of LOCKED DOORS undermine his rationale for the existence of "distant" epistemic luck. Moreover, distinguishing such a new type of epistemic luck is pointless given that there is a simpler and

less revisionary explanation of the kind of epistemic luck instantiated in LOCKED DOORS and in the structurally equivalent specification of CLOCK available: in those cases, the relevant knowledge-undermining possibility is not the close possibility of believing p falsely when one actually believes p truly, but the close possibility of believing closely related false propositions.

To sum up the points made in this essay, I have shown that safety and sensitivity can be easily fixed in order to rule out CLOCK as a case of knowledge. I have then argued that when it comes to giving counterexamples to modal principles like safety or sensitivity, it is very important that the way the cases are described leaves no room for ambiguous interpretations of what the relevant error possibilities are. I have illustrated the point by arguing that ODOMETER is open to several interpretations, to the extent that once the structure of the case is made clear, the problem for safety and sensitivity simply dissolves. My last point has been that Freitag provides no convincing reason for the existence of "distant" epistemic luck. In Freitag's cases, the epistemic luck in play can be explained by the easy possibility of believing closely related propositions falsely.[9]

[9]Even if there is no distant epistemic luck, CLOCK does pose a challenge to Pritchard's definition of veritic luck, according to which a true belief is lucky just in case in most close possible worlds in which the agent continues to believe the target proposition and which the conditions for belief formation are the same as in the actual world, the belief is false (see Pritchard 2007, p. 280). Pritchard could argue that his definition can be easily amended by including a clause on closely related false propositions, but that would be still insufficient to explain other cases of veritic luck such as versions of Russell's clock case in which the stopped clock can be only consulted at 3:00, which is the time the clock stopped (suppose that, unbeknownst to one, one is in a room with sixty security doors and the only one that is unlocked gives access to the stopped clock reading "3:00" only at 3:00). In my opinion, this and analogous cases strongly suggest that epistemic luck should not be defined in modal terms. In fact, a widespread intuition in the literature on the concept of luck (as well as in the moral luck literature) is that luck, in general, is a matter of lacking control over the lucky event (see Broncano-Berrocal (2015) for a positive account of luck in terms of lack of control). Surprisingly, the lack of control intuition is not taken into consideration when it comes to cases of epistemic luck. It is beyond the scope of this paper to explain how the notion of control can be used to account for the notion of epistemic luck and how that bears on the concept of knowledge.

Acknowledgements

Many thanks to Chris Kelp and two anonymous reviewers.

References

Broncano-Berrocal, F. (2014) "Is Safety in Danger?" *Philosophia* 42: 63-81.

Broncano-Berrocal, F. (2015) "Luck as Risk and the Lack of Control Account of Luck." *Metaphilosophy* 46: 1-25.

Freitag, W. (2014) "Safety, Sensitivity and "Distant" Epistemic Luck." *Theoria* 80: 44-61.

Goldman, A. (1987). "Nozick on Knowledge: Finding the Right Connection". In *The Possibility of Knowledge: Nozick and His Critics*, edited by S. Luper-Foy. Totowa, NJ: Rowman & Littlefield: 182–196.

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

Graham, P. (*forthcoming*) "Functions, Warrant, History". In *Naturalizing Epistemic Virtue*, edited by A. Fairweather and O, Flanagan. Cambridge: Cambridge University Press: 15-35.

Greco, J. (2000) *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press

Manley, D. (2007) "Safety, Content, Apriority, Self-Knowledge." Journal of Philosophy 104: 403–23.

Nozick, R. (1981) *Philosophical Explanations*. Oxford: Clarendon Press.

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

Pritchard, D. (2007) "Anti-Luck Epistemology." *Synthese* 158: 277–297.

Pritchard, D. (2015) "Anti-luck Epistemology and the Gettier Problem." *Philosophical Studies* 172: 93-111.

Russell, B. (1948) *Human Knowledge. Its Scope and Limits*. London: George Allen and Unwin.

Sosa, E. (1991) *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.