Swarthmore College

Works

Philosophy Faculty Works

Philosophy

11-1-2014

No Luck With Knowledge? On A Dogma Of Epistemology

Peter Baumann

Swarthmore College, pbauman1@swarthmore.edu

Follow this and additional works at: <https://works.swarthmore.edu/fac-philosophy>

Part of the [Philosophy Commons](#)

[Let us know how access to these works benefits you](#)

Recommended Citation

Peter Baumann. (2014). "No Luck With Knowledge? On A Dogma Of Epistemology". *Philosophy And Phenomenological Research*. Volume 89, Issue 3. 523-551. DOI: 10.1111/j.1933-1592.2012.00622.x
<https://works.swarthmore.edu/fac-philosophy/339>

This work is brought to you for free by Swarthmore College Libraries' Works. It has been accepted for inclusion in Philosophy Faculty Works by an authorized administrator of Works. For more information, please contact myworks@swarthmore.edu.

No Luck With Knowledge?

On a Dogma of Epistemology¹

Peter Baumann

Abstract

Current epistemological orthodoxy has it that knowledge is incompatible with luck. More precisely: Knowledge is incompatible with epistemic luck (of a certain, interesting kind). This is often treated as a truism which is not even in need of argumentative support. In this paper, I argue that there is lucky knowledge. In the first part, I use an intuitive and not very developed notion of luck to show that there are cases of knowledge which are “lucky” in that sense. In the second part, I look at philosophical conceptions of luck (modal and probabilistic ones) and come to the conclusion that knowledge can be lucky in those senses, too. It also turns out that a probabilistic notion of luck can help us see in what ways a particular piece of knowledge or belief can be lucky or not lucky.

The relation between knowledge and mere true belief which does not amount to knowledge has long kept philosophers busy. A number of authors hold that non-accidental true belief is knowledge (cf., e.g., Unger 1970, 114-115). Whatever one might think about that, there is widespread consensus that there cannot be any lucky knowledge. Or: Knowledge is incompatible with luck. More precisely, a belief which is true only by luck cannot

constitute knowledge. Let us call this the “No-Luck-Thesis” (see amongst many: Unger 1970, 114-115; see for one of the few dissenters Hetherington 2001, 84, and Hetherington 2011, ch.3):

(NLT) Necessarily, for all subjects *S* and propositions *p*: If *S* knows that *p*, then *S*’s belief that *p* is not true by luck.

This thesis needs, of course, to be made more precise. First: What does it mean for a belief to be “true by” luck? And: What is meant by “luck” here? Before we go into these questions, some preliminary clarifications are necessary. There is a general notion of luck with seems to apply primarily to events and states.² A state or event is lucky not in itself but for someone: in the case of good luck it is good for it, in the case of bad luck bad.³ Some states or events are neither good nor bad for someone, they lack evaluative significance for that being. Unexpected rain might constitute good luck for farmers, bad luck for picnickers, and neither good nor bad luck for distant observers. Apart from evaluative significance, luck also requires that the relevant event or state was not to be expected, is surprising, could easily not have happened or come about, or was improbable (much more on this below; see for the general conditions of luck: Rescher 1995, 32, *passim*; Riggs 1998; Pritchard 2005, ch.5; Coffman 2007).

People sometimes say that luck involves the lack of control (see on this: Rescher 1990; Katzer 1996, 106-111; Vahid 2001; Pritchard 2005, 127; Riggs 2007; Coffman 2007, Lackey 2008, 256-260). If “lack of control” means total or nearly total lack of

control, then we certainly have to say that lack of control is not sufficient for luck: That I don't have any control over the tides does not make me lucky in that respect. Lack of control in this sense is also not necessary for luck because many things are to some degree but not completely under my control and still lucky, like my lucky shot in a basketball game. If, by contrast, we mean lack of total or almost total control by "lack of control", then we also don't get quite the right connection between luck and control: That I am a brilliant but still not perfect basketball player does not mean that all my shots are lucky (but cf. Engel 1992, 60-64 for the idea that limited control entails luck in epistemic cases). However some degree of control (something between total lack of control and total control) seems necessary for luck: A lucky basketball shot requires a certain degree of skill but not too much. But how much control and how much lack of control is necessary or sufficient for luck? It is hard to see how one should answer these questions in a principled way. Apart from that, the notion of control relevant here is either a probabilistic or a modal notion (see below; for a more negative evaluation of the control thesis see Lackey 2008, 256-260). Hence, we should turn to straightforwardly modal or probabilistic accounts of luck. I think the best explanation of all this can be given in probabilistic terms: If an event is lucky, then the probability that it would happen, given certain conditions, is low enough or below a certain value (but cf. Morillo 1984; see also below).

In discussions of the relation between knowledge and luck a broad notion of luck should be used. Whether one thinks that having true beliefs is always or usually a good thing or not does not matter here. What matters is the accidentality or non-accidentality of a true belief, not its evaluative significance (see below for more details on

“accidentality”).⁴ The question is whether knowledge is compatible with accidentally true belief, and for that it does not matter what the evaluative significance of true beliefs is (but see Ballantyne 2011 and forthcoming). This is an advantage because it means that we can leave the complicated issue of epistemic values aside here. In the following I will use the term “luck” in this broader sense of accidentality; this is also the way it is used in the literature (though the difference between the broader and the more specific notion of luck is not often noticed or mentioned).

Furthermore, true belief can be called “lucky” in several different ways (see Engel 1992, Harper 1996, Pritchard 2005, chs. 5, 6).⁵ (NLT) is not specific enough as it stands: Some kinds of luck are uncontroversially compatible with knowledge whereas others aren’t. A subject could be lucky to have certain rare abilities (ability luck) which alone enables her to acquire certain kinds of true belief and knowledge (a mathematician might be lucky to be talented enough to prove that Fermat was right). Or a subject could be lucky to find herself in circumstances (circumstantial luck) which enable her to gain true beliefs and knowledge of a certain kind (e.g., a witness of a crime who just happens to be at the scene of the crime when it happens). These kinds of luck are unproblematically compatible with knowledge. What matters here is a different kind of luck. Following Thomas Nagel’s classification of types of moral luck (see Nagel 1979 and also Williams 1981), one might want to call this “resultant epistemic luck”. A true belief could be lucky in this sense even given the situation of the subject at the time of the acquisition of the belief (including their abilities, circumstances and available evidence), that is, the belief can be lucky even if the belief is not lucky in the other ways mentioned. Engel 1992 was one of the first authors to

clearly distinguish between different forms of epistemic luck; the one relevant to (NLT) he calls “veritic luck” which he explains as luck that a given belief is true, given the evidence for it (see Engel 1992, 67); in contrast, “evidential” luck is the luck of having the evidence in the first place (see Engel 1992, 67). Pritchard 2005 also uses the term “veritic luck” for the kind of luck relevant here and explains it as luck that a given belief is true (see Pritchard 2005, 146); he spells this out in terms of a safety account of luck and knowledge (see Pritchard 2005, 163; more on that below; see also Hetherington 2011, 90-98). As we will see below, there are several different ways of explaining what resultant epistemic luck is. (For more on epistemic luck, apart from Engel and Pritchard, see: Ravitch 1976; Foley 1984, 1987, ch.4; Latus 2000, Vahid 2001.) Whenever I speak of luck here, I have in mind some notion of resultant epistemic luck, that is, not evidential, circumstantial or “ability” luck.⁶ Whatever the details, we already know that we ought to modify (NLT):

(NLT*) Necessarily, for all subjects *S* and propositions *p*: If *S* knows that *p*, then *S*’s belief that *p* is not true by resultant epistemic luck.

In the following, I will use the term “luck” exclusively for “resultant epistemic luck”.

Is (NLT*) too strong? Does it not demand for knowledge that “the situation of the subject at the time of the acquisition of the belief (including their abilities, circumstances and available evidence)” (see above) guarantee that the belief be true? Isn’t (NLT*) “infallibilist” in the sense that it demands that the description of that epistemic situation entails the truth of that belief? Few people would accept such a strong condition. However,

(NLT*) does not entail any of that and is not committed to any strong claim like that. The absence of epistemic luck does not entail that a true belief is guaranteed. Consider a non-epistemic case. A very skilled archer hits a target that is not very difficult for him to hit. However, even for him there was a chance of failure and it was not impossible that he might miss the target; but that does not mean that he was lucky to hit it. Similarly in epistemic cases: I might, for instance, come to know something on the basis of excellent evidence. That the evidence does not entail the true proposition I come to believe on its basis and that there was a chance of acquiring a false belief does not mean that the belief was lucky. The notion of luck should thus not be understood and is usually not understood in the “infallibilist” sense that whenever there is even a slim chance or small possibility that the subject might acquire a false belief, the resultant true belief should count as lucky. It will become even clearer and more explicit below why and how (NLT*) is not committed to any such implausibly strong, “infallibilist” claim.⁷

I will argue here against the widely believed No-Luck-Thesis. I will start with a discussion of cases, using an intuitive, pre-theoretic (and hopefully non-deviant) grasp of what constitutes luck or accidentality as well as knowledge (I, II). The result of this discussion is that knowledge is compatible with luck (understood in a certain way). I would not be too disappointed if only some readers agree with me on all or most cases but I hope that most readers will agree with me at least on some cases. My central thesis holds, strictly speaking, if only one of the following cases is convincing. It is also important to stress, though, that the result of the discussion of cases is only provisional. Semantic

“intuitions” concerning the correct use of a term like “luck” and “knowledge” can only have the first but not the last word (as Austin used to stress with respect to ordinary language in general). Theoretical reflection needs to complement intuitive judgments. Some more reflection on the notion of luck will show in what sense knowledge is and in what sense it is also not compatible with luck. In the second part of the paper I will first discuss what I take to be the most widespread conception of luck: a modal one (III). It turns out, again, that luck, thus understood, is compatible with knowledge. I will end (IV) with the explanation of a probabilistic conception of luck which has the advantage of helping us to understand why and in what sense knowledge is compatible with luck and why and in what sense it isn’t. It will turn out that our intuitive notion of luck needs to be clarified and reconstructed to a certain extent in order to be useful for judgments concerning the relation between knowledge and luck. All this will help to understand what it means for a belief to be “true by luck”. It will also throw light on the notion of knowledge.

I. Intuitive Cases

Consider Russell’s famous clock example (see Russell 1948, 98 as well as a brief hint in Hetherington 1998, 466): A wants to know what time it is, looks at a clock which indicates the correct time, say 12:15, and thus comes to believe that it is 12:15. Unbeknownst to A, however, the clock has stopped some time ago and, as a matter of sheer accident, just happened to indicate the correct time when A consulted it. A does not know the time, adherents of the No-Luck-Thesis would argue, because his belief was true only by luck.

This appears quite plausible. Consider the following follow-up scenario:

(a) A's watch has also stopped. After consulting the clock, A immediately sets her watch which from now on indicates the correct time. 2 minutes later, A consults her watch and comes to believe (correctly, again) that it is 12:17.⁸

Does A know this? Whoever denies knowledge in the original case will also tend to deny it in (a) because (a) is similar in all relevant respects. Again, A was just lucky in her belief turning out true.

Could one argue that there while there is luck in Russell's original case, there is no luck in (a)? I don't think so because in both cases we are dealing with a subject who is around a stopped clock and acquires a true belief about the time in close causal proximity to the presence of the stopped clock. However, if we change (a) slightly or change the description of (a) slightly, for instance by leaving out that the watch was set after a stopped clock and by stressing that it is a very reliable watch set after a clock indicating the right time, then we might well get the opposite result that the resultant belief was not lucky. Given the changed description, it was quite improbable that the subject would acquire a false belief; not easily could the subject have acquired a false belief. I will argue in more detail in sections III and IV below that an adequate notion of luck should incorporate such an element of relativity of luck. However, given the way (a) was described and

constructed, it is plausible to say that the subject is lucky in that case, too, not just in Russell's original case.⁹

Time passes:

(b) The next morning, someone, B, asks A what time it is. A looks at her watch and tells B correctly what time it is.

Do A and B know what time it is? The only interesting difference between (a) and (b) is that more time has passed. The mere passage of time, it would seem, cannot turn non-knowers into knowers or lucky beliefs into non-lucky ones (see Reed 2000, 66-67). Here is an analogy with non-epistemic luck. Suppose you were lucky enough to win the lottery. Now you're rich and you're lucky to be rich. Time passes and 10 years later you're still rich, not having benefited from any other sources of wealth. Even given that no later events or circumstances were threatening your wealth, you're still lucky to be rich. The passage of time does not make the luck go away. Also, it seems, if we denied knowledge in (a), then we should also deny it in (b).

Let us assume that A keeps consulting her watch for the next year. It still indicates the correct time (within the boundaries set by common sense – A has no use for atomic clocks):

One year after the clock incident – having consulted her watch many times -
A checks the time again and comes to truly believe that it is 13:19pm.

Now let us add some further details to the description of this case which are irrelevant to luck and see what happens:

(c) A owns an exquisite and expensive Swiss watch. It is the kind of watch which need not be set for decades. A does not wear it but keeps it, for safety reasons, in a drawer. The watch is very dear to her because it was a wedding gift. Quite often she takes it out of the drawer to show it to guests.

Whenever she does that, she consults it for the time. This kind of watch, she was told, is super-reliable and she has no reason to doubt it. At this moment, one year after the clock incident, having consulted her watch many times, she takes the watch out again and looks at it. It shows the correct time:

13:19pm. She comes to truly believe that it is 13:19pm now. Neither she nor anyone else has any clue that the watch was set after a “Russellian” clock, exactly one year ago. A has no doubt and truly comes to believe that it is 13:19pm now.¹⁰

Does she know that? My impression is that most of us would now grant A knowledge of the time. The original epistemic deficiency (lack of knowledge) in the clock incident has been removed. Or so it seems. However, no factors relevant to luck have changed between (a), (b) and (c). Hence, A would, I think, still count as lucky in (c). Hence, (c) is a case of lucky knowledge.

One could object that A is also not lucky in (c). If the watch were not giving the correct time but rather, say, a much later or earlier time, then A would notice that something is not right and that the watch does not give the correct time. Hence, the fact that A does not notice anything like that justifies A in the (more or less implicit) assumption that the watch is reliable. But under such circumstances, A is not lucky. As a reply, I would like to stress first that it is not crucial for the example that A acts on the information she gets from the watch or checks it (implicitly or explicitly) against other sources of information; she might even forget a few moments later what the watch indicated. She would still be lucky but also be a knower of the time. More important is another point. Even if we assume that A gets (implicit or explicit) confirmation that the watch is reliable, she might only get confirmation that the watch is not wildly off the beam but no confirmation that the watch is also reliable with respect to giving the exact time, say in minutes of the hour. But that is exactly what her belief is about. Without such more specific confirmation about the reliability of the watch in indicating the minutes of the hour and not just the rough time of the day, she is still lucky in her correct belief in the time but also a knower.

But how is that possible? Does time heal such epistemic wounds? How could the passage of time turn non-knowers into knowers? And how could it do that without turning a lucky belief into a non-lucky one? Perhaps I have overlooked further relevant factors? Perhaps the fact that A has consulted her watch many times is relevant? But consider this case as a control case for (c):

(d) A owns an exquisite and expensive Swiss watch. It is the kind of watch which need not be set for decades. A does not wear it but keeps it, for safety reasons, in a drawer. The watch is very dear to her because it was a wedding gift. Quite often she takes it out of the drawer to show it to guests.

Whenever she does that, she looks at it but curiously never consults it for the time. She admires the beauty of the watch so much that it never occurs to her to use it for such prosaic purposes. At this moment, one year after the clock incident, she takes the watch out again and looks at it. She wonders what time it is and for the first time consults it for the time. The watch shows the correct time: 13:19pm. She comes to truly believe that it is 13:19pm now. This kind of watch, she was told, is super-reliable and she has no reason to doubt it. Neither she nor anyone else has any clue that the watch was set after a “Russellian” clock, exactly one year ago. A has no doubt and truly comes to believe that it is 13:19pm now.

Does A know the time? One would have to run this question by a representative sample of competent speakers of English. Even though (d) is perhaps not as clear a case as (c), I still suspect that many people would attribute knowledge in (d). I certainly would. But how could a source of information which does not give us knowledge the first time give us knowledge simply because of repeated use (as in case (c); more on this below)? One thing seems clear: The original luck has stayed. Hence, we have a case of lucky knowledge.

Let us add one more variation of Russell's clock example. This one is close in some respects to (a) and reminiscent of Alvin [Goldman's \(1986, 45\)](#) thermometer example (see below):

(e) A is a watchmaker and has just finished another one of his famous handmade Fregeant-watches. To set the watch, the unsuspecting A looks at the "Russellian" clock and sets the watch after it. From now on, the watch indicates the correct time. These kinds of watches need not be set again for a long time. A then puts the watch into her display box containing 8 other watches of the same series. The other watches have been set earlier, after the same clock but when it was still running o.k. One hour later, B enters the store, wanting to buy a watch from the Fregeant series. B picks the one A had just put into the display. To B, they all look the same. B doesn't set the watch because A tells him that it has just been set. Later in the day, B gives the watch to C as a birthday present. The next morning, C consults the watch for the first time. Given the great reputation of Fregeant watches, C has no doubts and comes to truly believe that it is 7.30am. Neither C nor B, A or anyone else has any clue that the watch was set after a "Russellian" clock, many hours ago. So, C simply comes to believe that it is 7.30 now.

Again, does the subject know what time it is? I think, again, that the answer is positive. At the same time, it is pretty clear that the subject's true belief is lucky.

All these different variations on a Russellian theme suggest that there can be lucky knowledge. To be sure, we have to be careful with the notion of luck and I will say a bit more on different notions of luck below. But we can already see a good reason to suspect that the absence of luck is no necessary condition of knowledge. I find this first conclusion quite interesting. However, we have to take it as a provisional one, given that it is based on an intuitive and unanalyzed notion of luck, and will have to take a closer look at the notion of luck below.

How can one explain the compatibility of knowledge and luck? I think it is very plausible to assume that there are different kinds of parameters which are salient or not so salient in (a) but not in (c), (d) or (e). The different degrees of salience of these kinds of parameters can explain why we judge there to be knowledge in one case, but not in the other. There are many theoretical doors open here to explain this in more detail but I do not want or need to pursue this much further here. Suffice it to say that my bets would be on a contextualist account according to which the truth conditions of knowledge attributions can vary with the attributor's context (cf., e.g., [Cohen 1987, Lewis 1996, DeRose 1999](#); more on that below). According to some (Cohen, Lewis) salience of relevant parameters is one contextual factor which determines variable standards for knowledge. Here, I would mainly rely on the idea that salience drives our attributions of knowledge.

For instance, one could say that when judging (a) and (c) different factors are salient to the attributor: in (a) rather the questionable source of information when the watch was set and in (c) rather the reliability of the watch and the long history of (actually or potentially) giving the correct time. Or take (a) in comparison with (e). In the first case we

have a focus on the initial coincidence while in the second case (e) we focus much more on the reliability of the watch and its similarity to non-luck exemplars of the same kind and thus on the many other comparable and unproblematic watches. An overview of the different kinds of parameters which can be salient or not is beyond the aims of this paper. My main point here is rather the denial of (NLT*).¹¹

But shouldn't we expect that as knowledge appears as we move from (a) to (c) or (d) or even (e), so does luck disappear? No, I think this series of cases illustrates an important difference in the influence of certain parameters (like the passage of time) on our judgments about knowledge on the one hand and about luck on the other hand. What could explain why knowledge appears after time has passed while luck still does not go away? Here is a sketch which may be sufficient for the present purposes. Attributing knowledge to a subject involves the (more or less implicit) acknowledgement that the subject has certain cognitive abilities the use of which helps explain that the subject acquired knowledge. When we're thinking about a subject's knowledge or lack thereof, we're very much thinking about what the subject has been and is doing over time (including the use of more or less reliable instruments). This idea has affinities with certain agent-realist forms of virtue epistemology (cf., e.g., [Sosa 1992, 2007](#) and [Greco 2006](#)). In contrast, attributions of luck are less "track-record"-related and more "origin"-related: Whether some state or event is lucky is very much a question of how it originated (in a lucky way or not) while our attributions of knowledge are not that closely tied to questions of origin and more to matters like performance over time. Therefore, we're ready to "rehabilitate" the epistemic subject after a certain amount of time while we stick with our attributions of luck

even after time has passed. We are happy enough, for instance, to use the lucky subject as a source of information and call him a “knower” while retaining the initial judgment about luck.

II. More Cases

More examples similar to the ones above can be brought up. Consider inferences and the following argument for the incompatibility of luck and (inferential) knowledge. Suppose A (competently) infers q from p and comes to believe in q . However, A does not know p but only has a lucky true belief that p . One could then argue that A’s belief that q is also lucky and for that reason does not constitute knowledge.¹²

However, one should be a bit skeptical about such an argument for the incompatibility of luck and (inferential) knowledge. First, a side issue: Why should A’s belief that q be lucky, too? Perhaps because it has been inferred from the content of another lucky true belief? This would suggest some principle like the following:

(TL) If S’s true belief that p is lucky and if S (competently) infers q from p , then S’s true belief that q is also lucky (if it was acquired on the basis of that inference alone).

But is (TL) true? Is luck is always transferred across competent inference (leading from the belief in the premises to the belief in the conclusion)? Consider this example. I come to believe that Joe is singing in the other room. And I am right. However, it is

usually Jack that does the singing around here and I cannot distinguish Jack's from Joe's voice. So, my belief that Joe is singing in the other room is true but lucky. However, if I infer from this that someone is singing in the other room, then the belief in that conclusion is not lucky. There are other interesting examples. Take any lucky true belief and assume that the subject correctly infers some truth of mathematics or logic from it. Joe is singing, hence $2+2=4$. Or: Joe is singing, hence Joe is singing or Joe is not singing. Should we assume that belief in the conclusion is lucky in such cases? Or suppose that I am driving in my car and have no idea whether I'm above the speed limit of 35 mph. I look at the speedometer and it reads "72 mph". And I am really going at 72 mph. However, my speedometer usually does not get the speed right exactly but diverges by a few miles from the true number. Thus, my belief that I am going at 72 mph is true by luck. Suppose I infer that I am way above the speed limit. Is belief in the conclusion really true by luck, too?

I think these cases raise some doubt about (TL) but I also think that they are not decisive against (TL) or similar principles. With respect to the singing case one could reply that belief in the conclusion is as lucky as belief in the premise and only appears non-lucky because there is a different inferential path which the subject could have taken instead of the actual one, namely one leading from a non-lucky belief in a different premise ("Jack or Joe is singing") to the (now) non-lucky belief in the conclusion that someone is singing. However, in the example above I did not take that path; hence, my belief in the conclusion only falsely appears to be non-lucky. So, this is a case of a lucky belief leading to another lucky belief which might falsely appear to be non-lucky. The other alleged counterexamples against (TL) are cases where belief in the conclusion can plausibly be

taken to be non-lucky but there is, in addition, another non-lucky belief in another premise in the background doing relevant inferential work. In the cases of logical and mathematical truths it is hard to imagine how we could not also, in addition, believe the conclusion on the basis of other premises which we believe in a non-lucky way. Hence, our inference is overdetermined and involves non-lucky beliefs in some other premises. The speedometer case can be dealt with among similar lines: How can I not notice that I am going way too fast by looking outside the window (which I am presumably doing when driving my car)?

So, even though there is, perhaps, some doubt about (TL) it is still pretty plausible. It might not hold in all cases and perhaps one will have to modify it. But there are many cases where luck transfers across competent inference from belief in the premises to belief in the conclusion. This is all we need here. Even if luck (and also the absence of luck) transfers across competent inference (always or in a given case), lack of knowledge doesn't (no matter whether knowledge does; see Baumann 2011).¹³ This has to do with the point mentioned above that the notion of luck is an origin-related notion whereas the notion of knowledge is a track-record-related notion. All this makes room for the possibility of lucky inferential knowledge - where luck transfers across competent inference while lack of knowledge does not.¹⁴

Here is a case where both the belief in the premises and the belief in the conclusion are lucky while the latter but not the former constitute knowledge:

- (f) Suppose that p and q are truths in number theory. Suppose further that A only has lucky true beliefs in p and in q and does not know them. Still, A

starts to use p as well as q as premises for further inferences. The chain of inferences gets longer and more complicated. After a large number of (correct) steps and a lot of time thinking things through, A reaches an extremely interesting conclusion C. A has managed to prove an important theorem in number theory. In the meantime, A has forgotten about p and q even though the whole proof rests on them. A still truly believes the conclusion C.¹⁵

Should we then still say that A does not know the conclusion of his proof (and only knows what the conclusion is)? This is implausible. However, whether or not luck is always transmitted across inference, it seems in this case that the belief in the theorem is as lucky as belief in p or in q on which the whole chain of inferences rest: Such inferences are shaky and unreliable. Again, it seems that we have to accept that there can be lucky knowledge. One might object that there are so many non-lucky steps in the chain of inferences that the subject is much less lucky (lucky to a much lesser degree) with respect to the conclusion than with respect to the premises p and q . Even if this is the case, it seems that the lucky reliance on the premises would still make the conclusion “lucky enough”.

Here is another variation:

(g) Detective D makes a competent and correct inference from a great number of premises all of which she knows (that the deed was done after

darkness, in the kitchen, with a knife, etc.) with the exception of one (only the psychoanalyst owns these kinds of knives) which she only believes (truly) by luck. Nonetheless, D comes to truly believe that the shrink did it.

D is lucky in her true belief but it seems very counter-intuitive that such a bit of luck can destroy knowledge of the conclusion.

One might object that a single lucky true belief in one of the many premises of an inference does not need to turn the true belief in the conclusion into a lucky one. Haven't I myself raised at least some doubt about the idea above that luck always transfers across competent inference (TL)? This might be true but there are still many cases where luck transfers (see above). Take case (g) and add (if that helps to raise the plausibility of my point) that the detective is really extraordinarily able in noticing the evidence, in weighing it and putting it together into one coherent overall picture. Few others would have been able to make the relevant inference and arrive at the true conclusion. Given that our judgments about knowledge are focused on the ability and track-record of the subject (see above), we should say that the detective comes to know the conclusion. For the judgment about luck, however, the factor of origin is salient and decisive. Given that one belief in one of the many premises was lucky but without it there would not have been a good inference to the true conclusion, we have to say that the belief in the conclusion is also lucky (even if it constitutes knowledge). Does it matter how much luck there is in the different beliefs in the different premises? Does a lucky belief in one unimportant premise necessarily make the belief in the conclusion of the whole inference lucky (against TL)? I

don't need to rule the possibility of such cases out categorically and am willing to leave tricky questions about the notion of importance of a premise aside here. What matters here is that case (g) is different and constitutes a case where belief in the conclusion is lucky: We may stress that the fact that only the shrink owns the relevant kinds of knives is quite important for the solution of the case and cannot be pushed into the background as unimportant.¹⁶

One could also reply to all this by introducing a notion of degrees of luck here and modify the No-Luck-Thesis along the following lines:

(NLT**) Necessarily, for all subjects *S* and propositions *p*: If *S* knows that *p*, then *S*'s belief that *p* is not true by a degree of resultant epistemic luck greater than threshold value *t*.

There are several problems with such a proposal. First, it is doubtful whether one can measure degrees of luck in a non-arbitrary way: Should one, in the above example, just go by relative numbers of premises? Or also weigh the importance of the different premises? But how? Second, how would one determine the threshold value *t* in a non-circular way, that is, without saying something along the lines of "t has to be just high enough to allow for knowledge"? Finally, the move to (NLT**) drastically changes the nature of the thesis. I myself am fine with the idea that knowledge is compatible with luck but as the remarks below show I don't spell this out in terms of acceptable degrees of luck. The typical

defenders of the No-Luck-Thesis, however, won't want to have anything to do with the kind of "softening" of it that we see in (NLT**).¹⁷

Consider other kinds of knowledge and this version of Goldman's well-known thermometer example (see Goldman 1986, 45):

(h) A wants to take B's temperature and is lucky enough to pick the only reliable thermometer from a basket of 100 thermometers (99 of which are, unbeknownst to A, not working properly). A takes B's temperature and discovers that B has got a temperature of 104°F.

Goldman and almost everybody else would judge that A does not know but only has a lucky true belief that the temperature is 104°F.

But is that true? Aren't we invited to take as a given in (h) that A has picked a working thermometer and ignore the chance and possibility to pick a defect one? But if we do that, then we should rather say that A only suffers from the unproblematic circumstantial epistemic luck, not from the problematic (for knowledge) resultant epistemic luck.¹⁸ In other words, we should then rather say that A's temperature-belief is not lucky. We might even want to count A as knowing the temperature. I don't want to deny that one can indeed construct a case like this or describe our case in this way. The crucial point here is that nothing forces us to characterize A's method of belief-acquisition as "measuring the temperature by using a working thermometer"; we might as well characterize it as "measuring the temperature by randomly choosing and using one of the

thermometers in the basket”.¹⁹ If we do the latter, then A must be characterized as suffering indeed from resultant epistemic luck, the kind of luck relevant here. What counts as circumstantial or resultant luck is thus also dependent on and relative to the description of the relevant case (here (h)). There is no fact of the matter that would determine whether a given case is a case of this or that type of luck (for more on this see section IV below). So, in (h) we do have a lack of knowledge and a presence of luck.

What if we imagine the following continuation of the case?

(i) A is so concerned about B that he decides to get B to the hospital where they do further tests over the next few days and finally discover that B is suffering from a rare viral infection. It is still early enough to treat B successfully. 5 days later, B starts to undergo extensive treatments. Two months later, B is released from the hospital. Another two weeks later, B feels normal again. The next weekend, he throws a party for his friends. He tells them, again and again, that “had A not known about my temperature, I would be dead now!”

If we add this to the original story, the inclination to deny knowledge to A about B’s temperature weakens considerably; I think we would even attribute knowledge to A – even though the element of luck has not disappeared. What is the additional parameter which makes the difference to our judgments?²⁰ Whatever our answer will be here (e.g. salience of later important consequences of the lucky true belief or salience of later important

actions of the subject and other agents) it is plausible to accept this as another case of lucky knowledge, against Goldman and many others (see Greco 2006 for a nice set of cases illustrating analogous points about moral luck).

Before we move on, I should reiterate a point made at the beginning. That knowledge is compatible with luck does not mean that it can be lucky in every respect. Some closer focus on the notion of luck can help to see that. Let us start with modal notions of luck (III) and then move on to probabilistic notions (IV). The last section will also explain in what sense knowledge is compatible with luck and in what sense it isn't.

III. Luck and Safety

What is luck anyway? So far, I have used an intuitive, pre-theoretic notion of luck, more precisely: of resultant epistemic luck. Let us now look at more theoretical conceptions of luck.

Some hold that knowledge just is non-accidental (non-lucky) true belief (see Unger 1970, 114-115). This might be too strong: Couldn't someone compulsively (and thus non-accidentally?) and truly believe that *p*? Be it as it may be, we can say this: Accounts of knowledge which explain the difference between mere true belief and knowledge will often (but certainly not necessarily) also suggest an explanation of what distinguishes knowledge from accidentally true belief and of what constitutes luck or accidentality (e.g., a defeasibility account of knowledge could identify luck with the absence of undefeasible justification; see Lehrer / Paxson 1969). There will then be as many or almost as many

different conceptions of luck as accounts of knowledge; typically, one is the shadow of the other. Few authors, though, have bothered to give an explicit account of luck in terms of their theory of knowledge. Modal accounts are an exception; let us therefore look at them.

Modal accounts explain luck in terms of what happens in close possible world: in terms of what could have easily happened. A true belief is lucky in virtue of how things stand in close possible worlds: whether it could have easily failed to be true. There are currently two main versions of this view. According to “sensitivity” views (modeled after sensitivity accounts of knowledge; see Nozick 1981, 172-196) the following holds:

(Sensitivity) A true belief is lucky just in case the person holds the belief in (all, most or many) close worlds in which it is false or the person does not hold the belief in (all, most or many) close worlds in which the belief is true.

According to a “safety” view (modeled after safety accounts of knowledge; see Sosa 1999 and, with much explicit reference to epistemic luck, Pritchard 2005 and also 2007) the following holds:

(Safety) A true belief is lucky just in case the person’s belief is false in (all, most or many) close worlds in which the person holds the belief.²¹

Because counterfactuals don't contrapose, the two views are not equivalent. Further details would have to be added to both accounts (e.g., a relativization to ways of belief formation) but we can skip these complications here.

What is interesting, though, is the idea of closeness of possible worlds. This idea is interesting also because there simply does not seem to be a unique acceptable closeness metric. Consider the following three worlds: the actual world (ACTUAL), a possible world in which the same laws of nature hold as in the actual world but in which I am a brain in a vat (VAT), and a possible world in which we are not brains in vats but the laws of nature are different from those in the actual world (LAWS). When we are thinking about epistemological skepticism, we tend to think that (LAWS) is closer to (ACTUAL) than (VAT) is. When we are thinking about the basic features of the world, we rather tend to think that (VAT) is closer to (ACTUAL) than (LAWS) is. Both closeness rankings are equally legitimate but they lead to incompatible closeness rankings (because they rely on different kinds of parameters which determine closeness). Given different and equally legitimate closeness rankings, we might end up with different but equally legitimate judgments (if we adhere to modal notions of luck) about the presence or absence of luck. And hence, if we adhere to the No-Luck-Thesis, we may end up with different judgments about the presence or absence of knowledge. All this suggests a hidden relativity in the concept of luck (luck is relative to a given closeness ranking) or a hidden context-sensitivity ("luck" has application conditions which vary with the closeness ranking given by the speaker's context). Most if not all of those who favor modal accounts of luck, safety and knowledge (e.g., Nozick 1981, Sosa 1999, Williamson 2000, Pritchard 2005) would

resist this kind of move but I think considerations like this one give good support for contextualism (but I cannot go into this here).²²

Even more important here is the following point. Above (sections I and II), we came to the conclusion that there appears to be an element of luck in all the different versions of the different cases discussed above. But, interestingly, in some cases we attribute knowledge, in others we deny it. So, there can be lucky knowledge. This result is confirmed if we check modal accounts of knowledge and luck against cases. According to sensitivity accounts of knowledge, knowledge requires the absence of luck in the sense explained by (Sensitivity); according to safety accounts of knowledge, knowledge requires the absence of luck in the sense explained by (Safety). Here is a case of luck in the latter sense which would still qualify as knowledge (similar cases can be made against sensitivity accounts of knowledge; because of the generic similarity of the two accounts I won't go into that here and just restrict myself to safety accounts; for more on these see Baumann 2008):

(k) It is 5 past midnight on January 1, 2001. Jack just finished his first letter ever to his old friend Jill. Jack knows the time and date and comes to believe that he finished his first letter to Jill in the 21st Century (see Baumann forthcoming; for more examples and objections against the idea that safety is necessary for knowledge see Gundersen 2003, 118-119, Neta/Rohrbaugh 2004, 399-400, Roush 2005, 118-126 and Comesaña 2005, 397; see also Sosa 2003, 159).

It seems clear that Jack knows that he finished the first letter to Jill in the 21st Century. Does it matter that he also adheres to the very popular false belief that the 21st Century began on January 1, 2000? If this were to show that Jack does not know this, then many people, perhaps the majority, would turn out not knowing that they were in the 21st Century when it began. This, however, seems clearly false.

However, if we assume (plausibly) that there is a close world in which Jack finishes his letter 8 minutes earlier, then we also have to admit that Jack's true belief in the actual world is lucky in the sense of (Safety). Given his calendaric confusion, Jack would have the false belief that he is already in the 21st century when he finishes the letter 8 minutes earlier. Hence, if we adhere to a safety account of knowledge and luck, we still have to admit that there are cases of lucky knowledge (given at least some admissible closeness rankings for possible worlds; see above; for more detail on this kind of critique of the safety account see Baumann 2008). Similar things hold for sensitivity accounts. This is interesting because proponents of such modal accounts of knowledge typically adhere to the No-Luck-Thesis and see their accounts also as theoretical foundations of this thesis (cf., e.g., Pritchard 2005).²³

One might object to this conclusion and point out that the relevant modal accounts of knowledge could be further developed so as to be able to accommodate the above difficulties. I have my doubts about the prospects but cannot go into this here; it is sufficient for me here to show that the basic version of, say, a safety account allows for lucky knowledge.

Does the safety account correctly identify necessary or sufficient conditions of luck? Does it confirm our “intuitive” judgments about luck? The answer does not matter that much, given that we’re looking at theoretical reconstructions of the notion of luck here anyway where it is not that important whether they match with an ordinary notion of luck. Suffice it to point at cases here which at least *prima facie* suggest negative answers. Jack’s belief in the letter case above is not safe but, according to ordinary intuitions (though not according to the safety account), might not count as lucky and might count as a case of knowledge: this is unsafety combined with the absence of luck. Conversely, it seems that one can be both safe and lucky in the intuitive sense: Suppose I wonder who the winner of the national poetry contest will be. Unbeknownst to me, the jury always makes sure it will be the president. I take a lucky but safe guess that it will be the president. More could be said about such cases (see Lackey 2008, 260-266 as well as Lackey 2006 for a nice discussion along similar lines) from both sides but I won’t go into it here since, as pointed out before, complete congruence of our theoretic notion of luck with the ordinary one isn’t the goal here.

IV. Probabilistic Luck

Probabilistic accounts of luck do not rely on modal notions, which is an advantage in the eyes of some (presupposing a non-modal notion of probability, that is, a notion of probability which does not rely on modal logic but rather on the probability calculus). An event is lucky, one could say, just in case it is highly improbable that it would happen

(where “improbable” and other terms are not understood in the subjective sense). In the case of the acquisition of a true belief (T), this gives us the following explanation:

$$P(T) \leq a$$

for a threshold value a typically if not always < 1 .

This, however, is not very satisfying as it stands: It is not easy to make sense of unconditional probabilities like the probability that someone acquired a true belief about something. What, for instance is the probability of rain here and now? It is hard to make sense of this if we don’t indicate conditions: Given that it is August and we’re in Sicily, the probability is very low; given that it is December and we’re in Glasgow, the probability is much higher. If one is not convinced by this general skepticism about unconditional probabilities (but cf. Hájek 2007), then one can just restrict the discussion to the epistemic case. Is there an unconditional probability that I acquire a true belief about the presence of butter in the fridge? We should rather say that there is a probability that I acquire such true beliefs, given, say, that I look, etc. Only conditional probabilities seem to make sense here.

What we need then is the notion of conditional probability, also for a probabilistic notion of lucky true belief. A true belief is lucky just in case the probability that the person would acquire a true belief on the question whether p , given certain circumstances and ways of inquiry, pieces of evidence, etc., is below a threshold value a (again taking “probability” in a non-subjective sense here). For instance, Julie’s true belief that there is

an Air France airplane flying by is lucky insofar as the probability of getting this right, given that she just looks at the sky, is very low. At the same time, the probability might be quite high, given that she has just taken her farsight eyedrops. We can say then that

(Probability-Luck) A true belief is lucky just in case the conditional probability that

(T) the person acquires a true belief about whether p , given that she is in certain circumstances and uses certain ways of belief acquisition (C), is low or below (or at) a certain threshold value a : $(P(T/C) \leq a)$.²⁴

This notion of luck lends itself as much to relativism or contextualism as the modal notions mentioned above: C (as well as a – but I will focus here on C) varies with the (speaker's etc.) context. But in contrast to those modal notions, this conditionality is explicitly built into the account of luck here. This is a remarkable advantage of the probabilistic account of luck over the modal accounts.²⁵ Whether S is epistemically lucky varies with C. It is thus elliptical and potentially misleading to say of some event, state or true belief that it is lucky; rather, one should say that it is lucky with respect to some condition while it is not lucky with respect to some other conditions. There is no single, unique, C upon which one has to “conditionalize”.²⁶ There is also no such thing as *the* (unconditional) probability of acquiring a true belief. We can call this the “Relativity of Luck”. One further aspect of the relativity of lucks concerns the type of: Whether

something counts as circumstantial epistemic luck or as resultant epistemic luck (see above) depends on what we chose as our C. For instance, if in our case (h) above we take it as a condition (method of belief acquisition) that A chooses a working thermometer for his measurement, then he counts as only suffering from circumstantial luck. If, however, we take it as a condition (method of belief acquisition) that A chooses one of the many thermometers in the basket, then he counts as suffering from the kind of epistemic luck that is interesting here, namely resultant luck.

A lot of the current debate on epistemic luck, though, works with the implicit assumption that there is such a unique probability and that luck is non-relative or absolute.²⁷ Many discussions of cases used to support a particular account of luck usually suffer from reliance on this assumption of uniqueness and absoluteness: pre-theoretic intuitions about whether a case is a case of luck rely on implicit assumptions what one has to take as given; change what is given and the judgment about luck might change drastically. This explains the impression of indeterminacy one can get when looking at “the same case” in the light of different further assumptions. A relativist or contextualist account (like the probabilistic one above) fares much better here than absolutist accounts (e.g., the modal accounts mentioned above). A more detailed relativist or contextualist account of probabilistic luck would also have to explain why it is that we often seem to agree in our judgments about luck (e.g., in the original clock case). This, however, is clearly beyond the scope of this paper.

We can apply this probabilistic notion of epistemic luck to our initial case and its variations. Notice that the probabilistic account does not have a problem with clock cases.

A's true belief that it is 12:15 is lucky insofar as the probability of getting the time right was very low, given that A consulted a stopped clock. In a similar way, A is lucky in the other versions of the case, for instance in (a) (see above); similar things hold for the other cases (which I won't deal with in detail here). All this can change, though, if we change what we take as given; then our verdict might be "not lucky!". But, as pointed out above, we do think that in the later versions of the case A knows the time. In other words, using the probabilistic account of luck, we also get to the conclusion that knowledge can be lucky.

There is thus lucky knowledge. But what drives or should drive our judgments about knowledge? It is very plausible to pair (Probability-Luck) with

(Probability-Know) If a true belief is knowledge then the conditional probability that

(T) the person acquires a true belief about whether p ,
given that she is in certain circumstances and uses certain ways of belief acquisition (C), is high or above a certain threshold value a : $(P(T/C) > a)$.²⁸
(We might, perhaps, even strengthen this to a bi-conditional²⁹)

But don't (Probability-Luck) and (Probability-Know) simply entail that knowledge requires the absence of luck? Am I not shooting myself in the foot here? No. If we relativize in the same way and to the same conditions C (again, I will leave the variability of the threshold value aside here) when we make judgments about luck and judgments

about knowledge, then there cannot be lucky knowledge. The main argument of this paper would thus not go through. However, what forces us to tie the notions of luck and knowledge so closely together? Sure, if we assume that knowledge just is non-lucky true belief, then we have to coordinate our judgment in this way. But this would be begging the question in favour of (NLT*). Hence, we should be free to relativize to different conditions when we make judgments about luck and about knowledge. And we do in fact do that. Even if both knowledge and luck allow or even ask for a contextualist analysis (something I cannot go much into here), they can come apart because attributions of knowledge are sensitive to further factors, like salience of temporal distance or important consequences (see above) to which attributions of luck (which are more origin-related) do not seem to be so sensitive. Hence, the “C” (and the “a”) in (Probability-Luck) and in (Probability-Know) can differ such that lucky knowledge is possible. The examples in sections I and II above illustrate this possibility.

A probabilistic account of luck and knowledge thus offers its own reasons to deny (NLT*), in addition to the cases mentioned at the outset.³⁰ It also helps us better understand in what sense knowledge is compatible with luck: The conditions (and also the thresholds) we refer to or use in our judgments about luck can diverge (and often do diverge) from those we refer to or use in our judgments about knowledge. At the same time, we can now also better understand why and how knowledge is not compatible with luck “across the board”: As (Probability-Know) tells us, if S knows that p , then there are some conditions C (and thresholds a) such that S’s true belief is not lucky with respect to them.

V. Conclusion

One might say that I have “only” appealed to psychological facts concerning the use of words like “knowledge” and “luck”. These facts do not tell us much, the objection continues, about how we ought to use these terms. I think this objection misses the target. How we ought to use a term, the rules for the correct use of it, are determined by facts about speakers and how they do in fact use the term. No matter what the psychological mechanisms influencing the use of certain terms are, it is very hard to imagine – except in very special circumstances involving expertise and other factors – that a majority of speakers could be systematically wrong about the meaning of a term. The above judgments about knowledge and luck are representative ones.

One could, finally, object that given a more sophisticated notion of luck, it would become clear that knowledge is not compatible with luck in that sense. Well, as long as such an account is not presented, it is futile to speculate about the mere possibility of such an account. The No-Luck-Thesis is formulated in terms of a primitive and general notion of luck. Knowledge is, it turns out, compatible with luck in this sense. Several cases were presented above; if only one of them can stand up, then the No-Luck-Thesis is in serious trouble; I think all of them go through. If we spell out the notion of luck more explicitly and in terms of current epistemological theories, we come to the same result. This is all I have been arguing for here, against what most epistemologists currently believe:

Knowledge is compatible with luck.³¹³²

References

- Alston, William P. 1995, How to Think about Reliability, in: *Philosophical Topics* 23, 1-29.
- Ballantyne, Nathan 2011, Anti-Luck Epistemology, Pragmatic Encroachment, and True Belief, in: *Canadian Journal of Philosophy* 41, 485-504.
- Ballantyne, Nathan, Luck and Interests, in: *Synthese* forthcoming.
- Baumann, Peter 2008, Is Knowledge Safe?, in: *American Philosophical Quarterly* 45, 19-31.
- Baumann, Peter 2011, Epistemic Closure, in: Sven Bernecker / Duncan Pritchard (eds.), *The Routledge Companion to Epistemology*, London: Routledge, 597-208.
- Baumann, Peter forthcoming, Safety, Virtue, Scepticism: Remarks on Sosa.
- Comesaña, Juan 2005, Unsafe Knowledge, in: *Synthese* 146, 395-404.
- Coffman, E.J. 2007, Thinking about Luck, in: *Synthese* 158, 385-398.
- Cohen, Stewart 1987, Knowledge, Context, and Social Standards, in: *Synthese* 73, 3-26.
- DeRose, Keith 1999, Contextualism: An Explanation and Defense, in: John Greco/ Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*, Oxford: Blackwell 1999, 187-205.
- DeRose, Keith 2009, *The Case for Contextualism. Knowledge, Skepticism, and Context*, vol.1, Oxford: Oxford University Press.
- Engel, Mylan jr. 1992, Is Epistemic Luck Compatible with Knowledge?, in: *Southern Journal of Philosophy* 30 (2), 59-75.

- Feldman, Richard 1985, Reliability and Justification, in: *The Monist* 68, 159-174.
- Fetzer, James H. 1977, Reichenbach, Reference Classes, and Single Case Probabilities, in: *Synthese* 34, 185-217.
- Foley, Richard 1984 Epistemic luck and the purely epistemic, in: *American Philosophical Quarterly* 21, 113-124.
- Foley, Richard 1987, *The Theory of Epistemic Rationality*, Cambridge/MA: Harvard University Press; (esp.: ch 4: Epistemic Luck, 174-208)
- Gendler, Tamar Szabó/ Hawthorne, John 2005, The Real Guide to Fake Barns: A Catalogue of Gifts for your Epistemic Enemies, in: *Philosophical Studies* 124, 331-352.
- Goldman, Alvin I. 1986, *Epistemology and Cognition*, Cambridge/MA & London: Harvard University Press.
- Greco, John 2006, Virtue, Luck and the Pyrrhonian Problematic, in: *Philosophical Studies* 130, 9-34.
- Gundersen, Lars Bo 2003, *Dispositional Theories of Knowledge*, Aldershot & Burlington, VT: Ashgate.
- Hájek, Alan 2007, The Reference Class Problem is Your Problem too, in: *Synthese* 156, 563-585.
- Harper, William 1996, Knowledge and Luck, in: *Southern Journal of Philosophy* 34 (3), 273-283.
- Hawthorne, John 2000, Implicit Belief and Apriori Knowledge, in: *Southern Journal of Philosophy* 38, supplement (Spindel conference 1999; David Henderson & Terry

- Horgan (eds.), *The Role of the Empirical (and of the Apriori) in Epistemology*, 191-210.
- Hawthorne, John 2004, *Knowledge and Lotteries*, Oxford: Clarendon.
- Hetherington, Stephen 1998, *Actually Knowing*, in: *The Philosophical Quarterly* 48, 453-469.
- Hetherington, Stephen 2001, *Good Knowledge, Bad Knowledge. On Two Dogmas of Epistemology*, Oxford: Clarendon.
- Hetherington, Stephen 2011, *How to Know. A Practicalist Conception of Knowledge*, Malden/ MA: Wiley-Blackwell
- Kaplan, David 1989, *Demonstratives*, in: Joseph Almog/ John Perry/ Howard Wettstein (eds.), *Themes from Kaplan*, New York/ Oxford: Oxford University Press, 481-563.
- Katzer, Charlotte 1996, *When Is Knowledge a Matter of Luck?*, in: *Grazer Philosophische Studien* 51, 105-120.
- Klein, Peter D. 2008, *Useful False Beliefs*, in: Quentin Smith (ed.), *Epistemology. New Essays*, Oxford: Oxford University Press, 25-61
- Kripke, Saul A. 2011, *Nozick on Knowledge*, in: Saul A. Kripke, *Philosophical Troubles. Collected Papers, Volume 1*, Oxford: Oxford University Press, 162-224.
- Lackey, Jennifer 2006, *Pritchard's Epistemic Luck*, in: *Philosophical Quarterly* 56, 284-289.
- Lackey, Jennifer 2008, *What Luck Is Not*, in: *Australasian Journal of Philosophy* 86, 255-267.

- Latus, Andrew 2000, Moral and Epistemic Luck, in: *Journal of Philosophical Research* 25, 149-172
- Lehrer, Keith / Paxson, Thomas D. jr. 1969, Knowledge: Undefeated Justified True Belief, in: *The Journal of Philosophy* 66, 225-237.
- Lewis, David 1996, Elusive Knowledge, in: *Australasian Journal of Philosophy* 74, 549-567.
- Luzzi, Federico 2009, Counter-Closure, in: *Australasian Journal of Philosophy* 88, 673-683.
- Madison, B.J.C. 2011, Combating Anti Anti-Luck Epistemology, in: *Australasian Journal of Philosophy* 89, 47-58.
- Morillo, Carolyn R. 1984, Epistemic Luck, Naturalistic Epistemology and the Ecology of Knowledge, or, What the Frog Should Have Told Dretske, in: *Philosophical Studies* 46, 109-129.
- Nagel, Thomas 1979, Moral Luck, in: Thomas Nagel, *Mortal Questions*, Cambridge: Cambridge University Press, 24-38.
- Neta, Ram / Rohrbaugh, Guy 2004, Luminosity and the Safety of Knowledge, in: *Pacific Philosophical Quarterly* 85, 396-406.
- Nozick, Robert 1981, *Philosophical Explanations*, Cambridge/MA: Harvard University Press.
- Pritchard, Duncan 2004, Epistemic Luck, in: *Journal of Philosophical Research* 29, 193-222.
- Pritchard, Duncan 2005, *Epistemic Luck*, Oxford: Oxford University Press.

- Pritchard, Duncan 2007, Anti-Luck Epistemology, in: *Synthese* 158, 277-297.
- Ravitch, Harold 1976, Knowledge and the Principle of Luck, in: *Philosophical Studies* 30, 347-349.
- Reed, Baron 2000, Accidental Truth and Accidental Justification, in: *The Philosophical Quarterly* 50, 57-67.
- Rescher, Nicholas 1990, Luck, in: *Proceedings and Addresses of the American Philosophical Association* 64 (3), 5-19.
- Rescher, Nicholas 1995, *Luck. The Brilliant Radnomness of Everyday Life*, New York: Farrar Straus Giroux.
- Riggs, Wayne 1998, What Are the “Chances” of Being Justified?, in: *Monist* 81, 453-472
- Riggs, Wayne 2007, Why Epistemologists Are so down on Their Luck, in: *Synthese* 158, 329-344.
- Roland, Jeffrey/ Cogburn, Jon 2011, Anti-Luck Epistemologies and Necessary Truths, in: *Philosophia* 39, 547-561.
- Roush, Sherrilyn 2005, *Tracking Truth. Knowledge, Evidence, and Science*, Oxford: Clarendon.
- Russell, Bertrand 1948, *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.
- Sosa, Ernest 1992, Generic Reliabilism and Virtue Epistemology, in: *Philosophical Issues* 2 (Enrique Villanueva (ed.), *Rationality in Epistemology*, Ridgeview: Atascadero, CA), 79-92.

- Sosa, Ernest 1999, How to Defeat Opposition to Moore, in: *Philosophical Perspectives* 13, 141-154.
- Sosa, Ernest 2003, The Place of Truth in Epistemology, in: Michael DePaul/ Linda Zagzebski (eds.), *Intellectual Virtue. Perspectives from Ethics and Epistemology*, Oxford: Clarendon, 155-179.
- Sosa, Ernest 2007, *A Virtue Epistemology. Apt Belief and Reflective Knowledge*, volume I, Oxford: Oxford University Press.
- Unger, Peter 1970, An Analysis of Factual Knowledge, in: Michael Roth/ Leon Galis (eds.), *Knowing. Essays in the Analysis of Knowledge*, New York: Random House, 113-130.
- Vahid, Hamid 2001, Knowledge and Varieties of Epistemic Luck, in: *Dialectica* 55, 351-362.
- Warfield, Ted A. 2005, Knowledge from Falsehood, in: *Philosophical Perspectives* 19, 405-416.
- Williams, Bernard 1981, Moral Luck, in: Bernard Williams, *Moral Luck. Philosophical Papers 1973-1980*, Cambridge: Cambridge University Press, 20-39.
- Williamson, Timothy 2000, *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Yamada, Masahiro 2011, Getting It Right By Accident, in: *Philosophy and Phenomenological Research* 83, 72-105.

Notes

-
- 1 This article has been published in *Philosophy and Phenomenological Research* 89, 2014, 523-551.
The definitive version is available at
<http://onlinelibrary.wiley.com.proxy.swarthmore.edu/store/10.1111/j.1933-1592.2012.00622.x/asset/phpr622.pdf?v=1&t=i26yv1ni&s=c3583a10383b9be5c2a95dc6ce6f51db84c92c00>.
 - 2 One may also want to include facts here. Persons are lucky in a secondary way, namely insofar as certain kinds of things happen to them or they find themselves in certain kinds of circumstances; we call such persons “lucky” because of the nature of those events and circumstances, not because of traits of the person. This might seem controversial but I don’t have to argue for it here because the kind of luck relevant here only applies to (epistemic) events and states.
 - 3 It is not necessary here to speculate about what kinds of beings or entities are susceptible to luck and what determines whether something constitutes good or bad luck for it: the being’s views about their own good or, rather, some more “objective” factor.
 - 4 I am using “accidental” as a synonym of “lucky” here (see also Yamada 2011). Should this raise concerns, I’d be happy to replace all talk about something being accidental by talk about it being lucky; nothing substantial hinges on these verbal matters here.
 - 5 We are not dealing here with luck concerning other epistemic states or events like having good warrant for some proposition. One can also be lucky not to acquire a false belief, when acquisition of a false belief is to be expected under the circumstances (also: one can be lucky to suspend belief, when that is the right thing to do but not what is to be expected). These kinds of epistemic luck lie beyond the topic of this paper.
 - 6 One might doubt that one can draw this kind of distinction between different kinds of luck in any principled way (see below). Even if that is the case, it is – as the remarks below will show – no problem for the view defended here.

-
- 7 Thanks to an anonymous referee for raising this issue. The question is also not how much luck is compatible with knowledge but rather how much of a chance or possibility of acquiring a false belief is compatible with the resultant true belief not being lucky and thus compatible with the true belief constituting knowledge. It should go without saying that there is only a fuzzy line between luck and the absence of luck.
- 8 John Hawthorne uses this very example to a very similar conclusion (see 2000, 202-203). See also DeRose 2009, 17-18 who shortly mentions this case and agrees with Hawthorne; however, he doesn't discuss it in detail and neither does Hawthorne. Cf. also Hawthorne 2004, 69.
- 9 Thanks to an anonymous referee for bringing up this point.
- 10 In this cases as well as in cases (d) and (e) below we are not assuming that the watch stopped at any time before but only that the first time it's set it's set after a stopped clock indicating the correct time. Hence, there is nothing wrong with the assumption that the watch is extremely reliable. Thanks to Darrell Rowbottom for bringing this aspect of the cases to my attention.
- 11 One might also think that the path from (a) to (c), (d) or even (e) goes through a slippery slope. But why? I do not see where the relevant vagueness would be supposed to lie here.
- 12 An inferred true belief can be lucky in more than one way. It can, for instance, be lucky insofar as the inference was not a competent one. In such a case, however, it is uncontroversial that the inferred belief does not constitute knowledge – no matter how one stands on the No-Luck-Thesis. The above case of inferential luck is more interesting here.
- 13 The following principle of “reverse closure” is not correct:

(RC) If S does not know that p and infers q from p , then S does not thereby come to know that q .

On the contrary, one can come to know something on the basis of an inference from something one does not know because it is false, as the case of an inference that someone is singing from a false premise that Jack is singing suggests. Here is another, perhaps more convincing example involving

a false premise. Assume that some theories or theoretical systems approximate the truth without being true in the strict sense. Suppose that Newtonian mechanics is such a case. Someone could use Newtonian assumptions to arrive at a true conclusion about the rough date of the next lunar eclipse. We would then, it seems, still credit the person with knowledge of the conclusion. See Warfield 2005 and Klein 2008 but also Luzzi 2009.

- 14 Thanks to an anonymous referee whose comments convinced me that one should be more confident about (TL) than I was initially.
- 15 If one thinks that a mathematical proof does not lead to knowledge (or even true belief) in the conclusion, then one can easily change the example without loss of substance. – For a related but still significantly different case see Hetherington 2011, 91.
- 16 Thanks to an anonymous referee who raised questions about (g) and provoked this defense of my characterization of the case as one of lucky knowledge of the conclusion.
- 17 Thanks to Darrell Rowbottom here.
- 18 Thanks to an anonymous referee for this objection.
- 19 See also fn.23 about the generality problem.
- 20 Sure, A might infer from the fact that B has a viral infection that the temperature must have been high (perhaps even that it must have been 104°F). However, our example also works without this additional assumption. Even without this kind of inference and even before he finds out about the viral infection, we would still credit A with knowledge (in case (i)). Thanks to Darrell Rowbottom for pressing me on this point.
- 21 Here is Pritchard's 2005, 163 formulation of the safety condition for knowledge (or, *mutatis mutandis*, for the absence of luck): "For all agents, φ , if an agent knows a contingent proposition φ , then, in nearly all (if not all) nearby possible worlds in which she forms her belief about φ in the same way as she forms her belief in the actual world, that agent only believes that φ when φ is true."
- 22 I am not claiming that there is total indeterminacy of closeness rankings. All I am saying is that there is enough indeterminacy to cause serious problems for modal accounts.

23 Apart from all that, the safety account (as the sensitivity account) has problems with some classic examples of Gettierization, for instance with Russell's clock example. How could A's belief that "it is 12:15 now" be false? Sure, the clock could have stopped at a different time (N:N) but then A would not have acquired the original belief "that it is 12:15 now" but rather the belief "that it is N:N now", with "N:N" for another time. These cases are, obviously, not relevant for the safety account because the belief has to be held fixed across variation of possible worlds. What, however, if A had looked at the clock at a different time, say at 12.20? Sure, he would have ended up with the belief that "it is 12:15 now". But this is a different belief from the original one, simply because of the indexicality of "now". A's utterance or thought at 12:15 that "it is 12:15 now" is different in content from A's utterance or thought at a different time N:N "that it is 12:15 now". All the relevant close worlds, however, are worlds in which the subject believes that it is 12:15 ("now"). The safety theorist has to hold fixed both the state of the clock and the time of the belief acquisition. This makes it very hard if not impossible to think of any "close worlds" in which the subject's belief is false. Still, the subject's belief is safe but she does not know what she believes. The safety theorist could concede this and reply that the safety condition was never meant as a sufficient condition of knowledge (only a necessary one). However, given our problem with the above kinds of beliefs, it would with this reply become quite unclear what relevance safety can have in the first case when one is dealing with indexical (or demonstrative) beliefs, beliefs which are simply trivially safe (cf. Roland/ Cogburn 2011 who briefly apply the safety view to the clock example but do, like many others, neglect the problem here).

It does not seem to help to take refuge to belief types rather than belief tokens where the notion of a belief type could be defined, say, in terms of David Kaplan's (1989) difference between character and content: Speakers who believe what they say when they utter the sentence "It is 12:15 now" share the "character" of their belief even though the contents of their beliefs (and thus their beliefs) vary with circumstances (we can ignore the details of such an account here). This kind of move from belief tokens to belief types or from content to character would drastically change the

safety account which is set up in terms of belief tokens. Many other beliefs of the “same character”, for instance, would have to be included in the safety condition. This seems way too inclusive.

Suppose I come to know that it is raining here now by looking out of the window and noticing that it is clearly raining here now. My belief token could not have been easily false but it is well conceivable that many other tokens of a belief of the same type could have easily been false because I am not good at detecting rain in different circumstances and distinguishing it from, say, snow in different circumstances (times and places). I assume the safety theorist would not want to go with the belief type (character) here and, implausibly, say that I don’t know that it is raining here now. Or suppose that I don’t know that it is raining here now even though I come to believe it: In this particular case, I went with (unreliable) auditory evidence while usually I take a look. Almost all belief tokens of the same type (character) would not have been easily false, given that I am a great rain-detector. I assume the safety theorist would, again, not want to go with the belief type (character) and, implausibly, say that I do know that it is raining here now. Similar problems arise with indexical versions of Goldman’s fake barn cases (“There is a barn here”). See also Gendler/Hawthorne 2005, 333-334 for a similar thought, as well as Kripke 2011, 169-171 (esp. fn.20, 23), 192 (fn.56), 213-214.

- 24 For simplicity’s sake I am leaving the fuzziness of the threshold aside here. - Reference to ways of belief acquisition raises the “generality problem” (cf., e.g., Feldman 1985, Alston 1995). I cannot and need not go into this here (as the following remarks will show implicitly). – By the way, this account of luck also lends itself to an explanation of the ways in which luck involves or does not involve the lack of control (see above). – Finally, it does not matter here whether one takes (Probability-Luck) as a final definition of some notion of luck or is open to further modifications; the crucial point here is unaffected by such issues.
- 25 A full-blown and detailed general comparison between modal and probabilistic accounts would be beyond the scope of this paper.

-
- 26 If there was such a unique way to conditionalize, then there would be a solution of the notorious reference class problem as applied to the problem of epistemic luck. However, there is very little reason to expect any such solution (cf., e.g., Fetzer 1977 and Hájek 2007). I cannot go into further details here.
- 27 cf., for instance, Pritchard 2005, 146 and Engel 1992, 67. To be sure, both authors in a way “relativize” luck: Pritchard to ways of belief formation with respect to which a given belief is safe or not safe or lucky or non-lucky (see Pritchard 2005, 163) and Engel to evidence with respect to which a given belief is lucky or non-lucky. However, they are both absolutist about luck insofar as they hold that there is always only one acceptable way to relativize (to ways of belief formation or to pieces of evidence).
- 28 This constitutes a fallibilist view of knowledge. – Again, I am passing over the fuzzy nature of the threshold here.
- 29 In that case, however, there would a problem with the knowability of lottery-proposition. If the right-hand side formulates a sufficient condition for knowledge, then one has to accept that one can know lottery-propositions – which many people would deny. However, this problem does not arise here because (Probability-Know) only formulates (in a rough way) a necessary condition of knowledge.
- 30 Pritchard 2004, 195 argues that probabilistic accounts of luck (and knowledge) are incompatible with causal determinism and thus not fully general in scope. This, however, is a misunderstanding, given that we rely upon the notion of conditional probability in the way indicated above. Each concrete event can be both causally determined and fall under an event-type which is correlated only probabilistically with certain other even-types. Even if my acquisition of a certain token of the belief that there is an Air France airplane passing by right now is causally determined, the probability of acquiring such a true belief, given that I am just looking at the sky, will typically equal neither 0 nor 1.

-
- 31 Hetherington 1998 (esp. 455-9, 463) argues in a very different way for the possibility of lucky knowledge: Knowledge can be “lucky” or “fallible” insofar as in some cases a knowing subject could have easily not come to know. However, that one almost lacked knowledge does not entail that one does not in fact have knowledge (to think otherwise is to commit the “epistemic counterfactuals fallacy”: 456). This point is well taken but not really uncontroversial and irrelevant to the arguments above. For all this see also Madison 2011. Hetherington also hints at the point that the cases we’re judging are not really spelled out enough (see above) but he does not go into much detail or analysis here (see, e.g., 1998, 456-7). See also Hetherington 2011, ch.3.
- 32 I am grateful for comments by Sven Bernecker, Darrell Rowbottom, anonymous referees, an audience at the University of California, Irvine and the GPPC group for work in progress in epistemology.