“Absent Contrary Indication”: On a Pernicious Form of Epistemic Luck, and its Epistemic Agency Antidote

Alfonso Anaya¹

Received: 28 September 2020 / Accepted: 29 August 2021
© The Author(s) 2021

Abstract

It is widely accepted that knowledge is incompatible with the presence of non-neutralized defeaters. A common way of addressing this issue is to introduce a condition to the effect that there are no non-neutralized defeaters for the belief that p (i.e. a “no-defeaters condition”). I argue that meeting this condition leaves open a possibility for defeaters to squander our knowledge. The no-defeaters condition can be fortuitously met, and as a result it can be met luckily. I shall argue that this kind of luck is inconsistent with knowledge. In order to prevent this pernicious form of luck I introduce a “defeaters-responsiveness” condition, according to which subjects ought to be disposed to adequately address defeaters if they were to arise (even if they in fact do not arise).

1 Introduction

Epistemic irresponsibility and luck tend to sit awkwardly with knowledge.¹ It would not be far-fetched to say that in looking for knowledge we (partly) aim to eliminate or at least attenuate the epistemically negative impact that irresponsibility and luck can bring to our cognitive lives. This venerable idea can be traced back at least to Plato (2009) and has taken central stage in contemporary epistemology thanks to Gettier (1963).

It is relatively uncontroversial that carelessly formed beliefs cannot constitute knowledge, even if they turn out to be true and the subject has warrant in their favour. Similarly, there is wide agreement in thinking that certain forms of (epistemic) luck are incompatible with knowledge, even if the resulting beliefs turn out to be true and the subject has warrant in their favour.² In response, epistemologists

¹ As we shall see, this general claim should be qualified, for at least not all kinds of epistemic luck represent a problem for knowledge (see Unger 1968; Pritchard 2005).

² This might be regarded as one of the lasting lessons from the Gettier Problem (1963).

✉ Alfonso Anaya

¹ Faculty of Philosophy, University of Warsaw, Warsaw, Poland

have introduced responsibility and anti-luck requirements for knowledge (see Nozick, 1981; Sosa, 2009; BonJour, 1980; Kornblith, 1983; Pritchard, 2005).

In this paper I shall focus on attempts to safeguard putatively warranted beliefs in the market for knowledge from the irresponsibility produced by the presence of non-neutralized defeaters for those beliefs. I shall argue that the “No Defeaters Condition”, aimed to prevent one form of irresponsibility, falls short of addressing the epistemic threats brought about by defeaters.

In order to show the insufficiency of the standard No Defeaters Condition (NDC), I begin by exploring in detail a threat for knowledge that arises from a kind of epistemic luck that has received little attention in current epistemological discussion. This kind of epistemic luck is connected to a form of epistemic irresponsibility that arises when subjects are *unprepared* to adequately address potential epistemic defeaters for their beliefs. I shall argue that this kind of luck is incompatible with knowledge.

Having a clear understanding of this kind of irresponsibility (and how to prevent it) matters beyond providing us with an improved formulation of a necessary condition for knowledge (in the form of an improved formulation of a condition for epistemically responsible belief). It is also important for wider discussions regarding the kind of agency involved in knowledge acquisition – I argue that it sheds light on an often-neglected aspect of epistemic agency. On my diagnosis, part of the reason why the problematic phenomenon has escaped most epistemologists has to do with a relatively narrow conception of what one *does* (and what one *needs to do*) when one comes to know something. According to these narrow conceptions, epistemic agency is mostly limited to assessing one’s evidence and adjusting one’s doxastic states in response to those assessments (Sosa, 1985, 2015).³ I argue that an adequate understanding of the neglected phenomenon shows that responsible epistemic agents should, in addition, *be prepared* to react appropriately to attacks when their beliefs enter the market for knowledge – and this readiness for action is better understood as belonging to our epistemic agency.

In order to appreciate the phenomenon I want to address, it will be useful to begin by looking closely at the phenomenon of epistemic defeat and the way philosophers have addressed this threat in their epistemologies. The main claim to be defended here is that an adequate account of the phenomenon of epistemic defeat ought to introduce a defeaters-responsiveness condition for knowledge, instead of the standard “no-defeaters condition” often put forward by participants to the debate. Thus, in §2, I explore contemporary epistemological discussions where it is regarded as relatively uncontroversial that knowledge is incompatible with the presence of (non-neutralized) defeaters. Often, the explanation put forward appeals to the ideas that it is epistemically irresponsible to hold a belief in the face of non-neutralized defeaters, and that knowledge is incompatible with irresponsible belief. In §3 I argue, first, that the no-defeaters condition can be luckily met; second, I argue that the kind of luck involved in this case is incompatible with knowledge. §4 advances my solution

³ Some sceptics about epistemic agency even challenge that this is a genuine form of agency (Engel 2009; Setiya 2013). For a robust account of epistemic agency in perceptual knowledge see Anaya (2021).

to this problem. I argue that a defeaters-responsiveness condition suffices to prevent the pernicious kind of epistemic luck I identify. On the proposal put forward, subjects must be *disposed to act* appropriately in the face of defeaters, even in cases where no defeaters are actually present. §5 briefly explores some consequences of my proposal in epistemology and philosophy of action.

2 Knowledge and Defeat

In current epistemological discussions there is wide agreement that an adequate account of knowledge should include a no-defeaters condition.⁴ Perhaps the most straightforward reason for this view is that there are many ordinary cases in which the presence of defeaters either makes one lose putative knowledge, or it prevents one from acquiring it. A natural thought seems to be that if defeaters threaten our knowledge, then an explicit requirement calling for the absence of defeaters should ensure that this particular obstacle for knowledge is no longer a problem.

For instance, I take myself to know that I was born on a given date on the basis that my parents told me so. But suppose one day I go through the family archive and I discover some documents which suggest that I was adopted and was born on a different date. Arguably, the presence of this information (or perhaps my coming to know it) make it so that I no longer know (perhaps I never did) that I was born on that day. Further details of this example should be spelled out, to determine whether or not I know that I was born on that day – it will depend in part on when I was actually born, on whether my parents' testimony was reliable (or whether I reasonably took it to be), on whether the documents I discovered were authentic, and so on. What is certain is that if I am to know that I was born on that day, then it either should be the case that no such defeaters exist, or I should effectively neutralize them, for instance by showing that the documents are fake or that, despite being authentic, they're wrong about my real date of birth.

2.1 The Varieties of Defeaters

For my purposes it is important to call attention to two prominent distinctions traced in the discussion on defeaters. The No-Defeaters Condition (to be fully characterized in §2.3) should be read as construing defeaters in line with the following definitions. On the one hand, the undercutting/rebutting distinction traces a distinction in the *way* a defeater speaks against a putatively warranted belief. An *undercutting* defeater is a consideration that suggests that S's belief that *p* was formed in a defective manner such that its warrant is called into question. A *rebutting* defeater is a consideration that speaks in favour of the falsity of *p* (Pollock & Cruz, 1999; Bergman 2005). On the other hand, the psychological/normative distinction traces

⁴ A prominent exception is Maria Lasonen-Aarnio (2010), see Smithies (forthcoming) for a response. In favor of the view we have, among many others, Sosa (1985, 2015), McDowell (1993, 1996, 2013), Bonjour (1980), Goldman (1979), Lackey (Lackey 2005), Goldberg (2018).

a distinction along the dimension of the psychological possession of the defeating considerations. A *psychological* defeater is a psychological state of the subject (commonly a belief) which, if true, would speak against a putatively warranted belief. A *normative* defeater is a non-psychological state of affairs, the obtaining of which, speaks against a putatively warranted belief, *and* of which the subject should be aware (Lackey 2005). A *normative* defeater obtains even if the subject does not believe that the state of affairs obtains, whereas a psychological defeater obtains even if the corresponding belief turns out to be false.⁵

Hilary Kornblith's (1983) physicist example can be used to vividly illustrate these distinctions. This ego-driven physicist's vanity is routinely exhibited by his failing to accept evidence that speaks against his views. Let us imagine that our physicist finds himself in an academic conference where a colleague of his presents him with counter-evidence for his theory. If the evidence suggested, for instance, that the experimental design in the physicist's work is flawed, then this evidence would constitute an undercutting defeater for his favoured theory. But if the colleague shared her own experimental results, which speak against the physicist's favoured theory, then the evidence would constitute a rebutting defeater for his belief in the theory. Kornblith's physicist, being an arrogant and dismissive person who routinely tunes out criticisms, would not form the belief that the evidence presented is damaging for his theory, and thus would not have a psychological defeater for his belief in the theory. Yet, it would still be the case that there are normative defeaters for his belief in the theory, for there are certain beliefs he *should have formed* which would undermine his belief in the relevant theory.⁶ But if the physicist, no doubt acting out of character, formed the belief that the evidence provided by the colleague is damaging to his theory, he would find himself with a psychological defeater for his belief (even if it turned out that the colleague's criticisms were wrong).

2.2 Why is Knowledge Incompatible With Defeat

The idea that the presence of defeaters (of any kind) can prevent one from gaining knowledge, or can make one lose putative knowledge, is not highly controversial.⁷ The examples used to illustrate the notion of defeaters might suffice to elicit

⁵ Recently, it has been argued that there is an additional kind of defeater, distinct from the ones discussed so far: higher-order defeaters. It remains unclear whether this is indeed a distinctive kind of defeat or whether it can be subsumed under the category of undercutting defeaters. For discussion see Christensen (2010) and Smithies (*forthcoming*). I leave this complication aside, for it is irrelevant for the proposal I put forward.

⁶ Here we run into delicate matters as it is not clear whether a normative defeater, to be so, needs to be a *true* consideration to the effect that the relevant belief might be unwarranted. It could be argued that only true considerations are normative defeaters, whereas misleading or false considerations are merely *seeming* defeaters. Further, it could be argued that whereas the former have the power of hindering our warrant, the latter lack such power, even though both kinds of defeater can be such that they make one's dismissal of them epistemically irresponsible. For the purposes of this paper, it will make no difference what position we take on these delicate issues.

⁷ See Lasonen-Aarnio (2010) for dissent. Although it is important to note that her arguments merely target defeaters which turn out to be misleading. Which suggests that she would agree that non-misleading defeaters *are* incompatible with knowledge.

the intuition that one cannot know that p if there are defeaters for one's belief that p . Yet, it is still worth exploring *why* the presence of a defeater for p is inconsistent with knowing that p . A natural explanation suggests that the presence of defeaters render our targeted beliefs epistemically irresponsible, and this, in turn, is incompatible with knowledge.

An advantage of this suggestion is that it allows us to explain why the presence of defeaters is inconsistent with knowledge, even in cases where ignoring defeaters could be truth-conducive. In the absence of this explicit link between the presence of defeaters and epistemic irresponsibility it becomes puzzling why certain beliefs should not be considered knowledge, given that they are true and are based on good grounds. Consider the following case: your friend has actually landed the CEO position in a famous company and she calls you to inform you about her feat. Later that day you read in a reputable news outlet that someone else got the job – as a matter of fact the outlet has made a mistake. You could ignore the outlet's information and continue believing that your friend got the job. Notice that if you did, you would have a true belief based on good grounds. But the intuition that the resulting belief would not amount to knowledge remains. This intuition is straightforwardly accommodated by a deontological approach, according to which a requirement for one's knowing that p is that one forms and sustains the belief that p in an epistemically responsible manner. In slogan form, *knowledge requires doxastic responsibility*. Arguably, forming and/or sustaining a belief by ignoring a defeater against it would be epistemically irresponsible. As the example illustrates, one could be epistemically irresponsible even when our belief turns out to be true and based on solid grounds.⁸ The relevant belief will not amount to knowledge unless we do something to address the underlying irresponsibility in our belief.

It is important to add a precisification to the general claim that the presence of defeaters render one's targeted belief irresponsible and, thus, unknowledgeable. For it is widely acknowledged that subjects might have in their possession reasons which serve to neutralize a defeater. For instance, suppose that in my CEO example, after reading the conflicting information in the news outlet, you check the company's website to corroborate that your friend effectively got the job. In such a situation, the misleading information in the news no longer poses a threat for your knowledge, for you can now determine that in all likelihood that the news outlet made a mistake – you have adequately neutralized the defeater. Thus, we should make our claim more precise to capture this observation: it is the presence of *non-neutralized*

⁸ This claim would be endorsed by Pritchard (2005, ch. 7) – for he claims that epistemic irresponsibility is the result (at least sometimes) of "reflective epistemic luck". A *true* belief exhibits this kind of luck when its being true is a matter of luck *from the perspective of the subject's epistemic situation* (i.e., on the basis of the subject's introspectively accessible evidence). On this account, a *true* belief can be epistemically irresponsible – for the irresponsibility of a belief stems from the fact that the belief is not appropriately supported by the evidence possessed by the subject. This is consistent with my claim that ignoring a defeater would be irresponsible even if the resulting belief is true. Below (§3.1) I explore in more detail how this account of epistemic irresponsibility interacts with the proposal I put forward.

defeaters against p that render a belief that p irresponsible – and for this reason, it is the presence of non-neutralized defeaters that is incompatible with knowledge.⁹

2.3 No-Defeater Conditions for Knowledge

Given that knowledge is incompatible with having non-neutralized defeaters, it might be thought that the most straightforward way to deal with this threat would be to posit a No Defeaters Condition (NDC) for knowledge. On this “brute force-strategy” (Lasonen-Aarnio, 2010) a belief that p will qualify as knowledge only if there are no defeaters for the subject’s belief that p .¹⁰ The correct reading of NDC, let us remember, construes defeaters as they were defined in §2.1, above.

If defeaters of any kind (i.e., psychological or normative) threaten our knowledge, then it would be sensible to construe NDC as encompassing both kinds of defeaters. Although this might be relatively clear it is important to spell out why a psychological reading of NDC would be insufficient to address the underlying problem. Notice that a requirement to the effect that the subject lacks any belief which would undermine or rebut her target belief¹¹ would leave unaddressed other forms of epistemic irresponsibility. It is important to highlight that the problem with this version of the no-defeaters condition is not its inadequacy but its insufficiency. Kornblith’s physicist illustrates how someone can meet the no-(psychological)-defeaters condition in an epistemically irresponsible manner. For this is a subject who secures satisfaction of the condition by ignoring damaging evidence, and thus fails to form the beliefs which would defeat her belief that p . As a result, mere absence of psychological defeaters turns out to be too weak a requirement, for it allows for a different kind of epistemic irresponsibility to take hold of one’s beliefs. The reason why Kornblith’s physicist is epistemically irresponsible (despite lacking any psychological defeaters for his belief) is that he ignores relevant evidence that he *should have taken into account* – i.e., there are normative defeaters for his belief (Goldberg, 2018).

At the very least, then, we need a stronger reading of the no-defeaters condition, one that requires that no psychological *or* normative defeaters are in place. On this stronger interpretation, a belief qualifies as knowledge only if the subject does not believe there to be any reason, *and* there is no reason she should be aware of, to think that her belief is incorrect or ill-founded. This condition would ensure that the presence of psychological or normative defeaters disqualify a belief from counting as knowledge. Interestingly, a consequence of this condition is that two subjects who are intrinsically the same could differ a lot in their epistemic standings as a result of the nature of the environment they find themselves in. For there is a clear sense in which the makeup of the epistemological landscape is beyond a subject’s control.

⁹ Fort the sake of brevity, in what follows I will use “defeater” to refer to “non-neutralized defeater” unless I explicitly state otherwise.

¹⁰ Among many others see Goldberg (2018) for an explicit endorsement of this condition, as well as Sosa (2015, 1985).

¹¹ For a useful summary of positions which endorse this interpretation of the no-defeaters condition (and critical discussion of it) see Bergmann (2000: 88–91).

Thus, two people might believe that p on the same basis, but one of them might fail to know that p merely because there is information she should be aware of which speaks against p . This shows that there is an important sense in which whether one is in a position to know depends on whether "the world makes the subject a favour" (McDowell, 1993).

Now, this general approach to dealing with the threat of defeaters, recall, has been characterized as employing a "brute force strategy". Certainly, this strategy seems to address the problem in a blunt way: by stipulating that that a necessary condition for knowing that p is that there are no (psychological or normative) defeaters for the belief that p . But let us remember that there is an independent justification for the introduction of NDC (beyond merely blocking problematic counter-examples): it serves the purpose of preventing a pernicious kind of epistemic irresponsibility of taking hold of our beliefs.

However, I shall argue that NDC fails to neutralize a further threat posed by epistemic defeaters. This threat is not entirely unconnected to epistemic irresponsibility, nevertheless the central form of the threat is that of a pernicious form of epistemic luck. NDC is a flawed condition because it allows for the presence of this pernicious form of epistemic luck.

3 Undefeated-Doxastic Luck and the Shortcomings of the No-defeaters Condition

Here I want to begin a critical analysis of the no-defeaters condition as stated above, by focusing on whether the phenomenon of epistemic defeat can threaten our knowledge in ways different to those widely recognized in contemporary discussions. I shall argue that such a further threat does exist and that incorporating NDC into our epistemology is insufficient to address it. It will become clear below that this threat is not completely unconnected to issues pertaining to epistemic irresponsibility, but it will take some time to spell out the exact way in which epistemic irresponsibility enters into the picture.

In order to provide a roadmap, let me briefly summarize my argument for thinking that there is such a further threat to our knowledge. I begin by noting that the No-Defeaters Condition, as it has been described so far, can be satisfied in a lucky fashion. This kind of luck, although different in important respects from the standard type of luck involved in Gettier cases, is incompatible with knowledge, or so I shall argue. Thus, my argument has two stages. First, I establish that a belief can luckily meet NDC (§3.1). Second, I argue that this kind of luck is incompatible with knowledge (§3.2). In developing the second part of my case, epistemic responsibility gets back into the picture: I argue that the kind of luck I identify is incompatible with knowledge because it leaves our beliefs vulnerable to an inadmissible amount of risk of being epistemically irresponsible.

Since satisfaction of the no-defeaters condition is compatible with the obtaining of this kind of luck, the condition fails to neutralize this threat to our knowledge. This will pave the way for my positive proposal: in order to know that p , it does not suffice that there are no defeaters for the belief that p ; in addition, the believer must

be appropriately sensitive to potential defeaters against p . But first things first. Let me begin by spelling out my negative case by fleshing out its first two stages. I leave an articulation of the positive proposal for §4.

3.1 Lucky Satisfaction of NDC

To show that NDC can be luckily satisfied, I shall describe a conceptually consistent hypothetical scenario where NDC is met luckily. Let us begin by appreciating that the no-defeaters condition *can be* met in a fortuitous manner, through no effort from the agent. To see this, note first that the no-defeaters condition merely states that if a subject's belief is to constitute knowledge, then there cannot be (non-neutralized) defeaters for her belief. As stated, the condition does not require that subjects make any contribution toward meeting the condition.¹² In particular, a subject might have no defeaters (of any kind) against her belief that p – but it could very well be the case that the absence of defeaters is in no way the product of the subject's active endeavours.

Consider the following illustration. Imagine a university security guard making her usual rounds after working hours at the psychology department. She passes by a lab with a curious display: a large white table with a single bright red apple atop. She forms the belief that there is a red apple on top of the table. As a matter of fact, there is nothing strange going on. In particular, there is in fact no information she should be aware of which could make her suspect that there is something misleading in the scene she sees, and she has no belief to the effect that something misleading is going on or that there is something wrong with her vision. In this case, the security guard's belief that there is a red apple on top of the table satisfies NDC. The guard has no belief which makes her think that her belief was formed in a defective manner, or which makes her doubt that there is a red apple on top of the table (i.e., she has no psychological defeaters for her belief, either undercutting or rebutting). In addition, there is no information *of which the guard should be aware*, which would make her think that there is no red apple on top of the table, or which would make her think that she formed her belief in a defective manner (i.e., there are no normative defeaters for her belief, either undercutting or rebutting). Nevertheless, we can easily imagine the situation in a way that the absence of defeaters obtains fortuitously – through no active effort from the agent. In this kind of case, the satisfaction of NDC is in no way the result of the subject's endeavours – i.e., it is not a feat attributable to the agent.

In this particular example, NDC is met because there are no defeaters of any kind: if there are no defeaters of any kind, then there are no non-neutralized defeaters. Thus, this is a case where NDC is met (fortuitously) because there are no defeaters for the relevant belief. But at least in some of these cases, *there could have been* such defeaters. For instance, suppose that our security guard is peering into a lab

¹² Of course, some epistemologists might defend this aspect of the condition, perhaps arguing that anything more would be overly demanding. I disagree. Allowing for the fortuitous satisfaction of the condition is not innocuous – I leave this part of my argument for §4.

equipped with a sophisticated lighting system which makes anything on the table look red. Although the system is off, it could have been on. Although no one lets the guard know that this system has been installed on the lab, someone could have let her know.

But my aim is to show, recall, that NDC can be met luckily. The mere fact that there are counterfactual situations where defeaters could have arisen does not suffice to show this. On a standard account of luck, an event/state is lucky when its occurrence/obtaining in the actual situation is accompanied by its failing to occur/obtain in some of the nearest (i.e., most similar) counterfactual situations.¹³ An illustrative example is that of being narrowly missed by a ball thrown at us by a skilled thrower. This event is lucky because although the ball does not hit me in the actual world, there is at least one (modally) near counterfactual situation where the ball hits me – i.e., very few things would have to be different for the ball to hit its target (e.g. my face). Here it is important to note an aspect about the example I have used to illustrate the notion of luck: the ball is thrown by a skilled thrower, which means that it must be a rare occurrence for her to miss her throw.¹⁴ If the thrower is really skilled then the explanation of why she failed must appeal, e.g., to an abnormal aspect of the situation, such as an abnormal wind (Sosa, 2010). It could be argued that once we construe the situation in this way, it is not so clear that the missed throw is a matter of luck: “of course she missed the throw, just look at the wind!” Importantly, the notion of luck I aim to illustrate does not require the presence of agents. Consider a situation where you are standing under a palm tree to shield yourself from the sun. Suppose that you move slightly to your left and a fraction of a second later a coconut falls on the spot you previously occupied. It is a matter of luck that the coconut did not land on your head.

Under this conception of luck, then, for NDC to be met luckily, there need be at least a nearby counterfactual situation where defeaters for the relevant belief do arise. Note that the obtaining of defeaters in modally distant conditions would not render the actual satisfaction of NDC lucky. For instance, going back to the security guard example, the fact that an evil demon could have manipulated the guard’s experiences to make it seem to her as if she was looking at a red apple when in fact there is none, *is not* the kind of situation that would render the actual satisfaction of NDC

¹³ Here I assume a modal account of luck as the one that can be found in Pritchard (2005). An alternative account of luck has it that an event is lucky when it has a low probability of obtaining (de Grefte 2020). I take it that the probabilistic account of luck poses no threat for the proposal I advance here – I assume that the examples I use would also be classified as lucky events by a probabilistic account of luck.

¹⁴ Since this lucky occurrence involves an agent, important consequences related to issues of attributability depend on the reason why such a skilled thrower failed in this situation. Consider Sosa’s (2010) treatment of competent performances. He would argue that such a miss from a skilled thrower must be the result of a defect in the thrower’s shape (perhaps she was drunk) or in her situation (perhaps it was uncharacteristically windy). For instance, if it was windy, then the blame for failing does not fall entirely on the agent. Or consider a case where the thrower uncharacteristically performs a defective shot, only to hit the target thanks to a gust of wind. Even if the shot is successful, Sosa argues, the success cannot be attributed to the agent. These considerations matter for Sosa because his project depends on identifying the conditions under which an agent can be credited with the success of particular performances.

lucky.¹⁵ For I assume that such a circumstance would be very distant modally from the actual circumstances – thus, it is not the case that easily the security guard could have been the victim of this kind of deception. Nevertheless, it is possible to imagine counterfactual situations very similar to the actual one, where the security guard does face a defeater. Let me illustrate. Our security guard, let us remember, spots a red apple atop a white table in a lab and forms a corresponding belief. Suppose further that, still in the (imagined) actual world, a psychologist leaving her office late sees the guard peering into the lab and decides to let her know that the lab is currently testing a sophisticated lighting system which makes anything placed on the table look red. Often, careless researchers leave the system on so she decides to check and let the guard know about this situation. She is on her way to inform the guard about this when she receives an unexpected phone call from her father's carer. Alarmed by the possibility that her father's frail health has gotten worse she takes the call, forgets about the lab and leaves the building in a hurry. Arguably, given this description of the actual world, there is a modally near counterfactual situation where a defeater for the guard's belief does arise. After all, given the details about the psychologist's quick exit, very easily could the guard have been given a reason to think that the apple might not be red. That is, there is a nearby counterfactual situation where she has a psychological defeater for her belief, i.e., the world where the psychologist receives the call slightly later and has plenty of time to inform the guard about the lighting system.

So far, we have an example where NDC is met, *and* where very easily there could have been a defeater for the relevant belief. Nevertheless, this still does not suffice to show that NDC can be met luckily. For an event/state to occur/obtain luckily, there has to be a nearby counterfactual situation where *it fails* to occur/obtain. But it is not entirely clear that the counterfactual situation I have described is one where NDC fails to be met. For we have to recall that NDC states that a belief *p* constitutes knowledge only if there are no *non-neutralized* defeaters for it. To show that NDC can be met luckily we need a case where NDC is met, but where a nearby counterfactual situation is such that the defeaters that arise for the belief remain *non-neutralized*.

As the security guard's case has been developed so far, it is unclear that the potential defeaters would remain non-neutralized. As the case stands, we have no reason to think that the guard would be unable to neutralize potential defeaters in the counterfactual situation. So far, my description of the case does not prevent the guard from walking into the room and taking the apple outside to inspect in under normal lighting. If she could do this, then NDC could still be met in the nearby counterfactual situation, which means that NDC would not be met luckily.

But in order to secure the case that I need, we only need to add a further stipulation to the case I have been developing. Imagine that the security guard I have been talking about has a profound contempt towards academics, which causes her to automatically dismiss any testimony provided by an academic. This means that

¹⁵ The consideration that such possibilities are very distant in the modal space ground Sosa's (2009) response to radical scepticism.

our guard would dismiss the testimony by the psychologist in the counterfactual situation I have described. Given this further stipulation, it is likely that the guard would continue believing that the apple is red in spite of the information given by the psychologist. In such a case, the guard would have a non-neutralized defeater for her belief that the apple is red. This means that, in the counterfactual situation, the guard's belief would fail to meet NDC, even though in the actual world her belief meets NDC. Since we have imagined the case in a way that the counterfactual situation is in the modal vicinity, we find ourselves with a case where NDC is met luckily – for very easily the corresponding belief could have failed to satisfy NDC.

Here, it is important to reflect on the impact that this last stipulation has in the case. In particular I want to highlight the fact that now it seems that our guard would behave irresponsibly in the counterfactual situation. For she would uphold her belief in the face of cogent testimony against it and would dismiss the latter for no good reason. The kind of luck I identify, then, is related to epistemic irresponsibility in the following indirect way: when a subject's belief that *p* satisfies NDC luckily, there is a modally nearby counterfactual situation where the subject's belief that *p* is irresponsibly held. The aim of this paper, recall, is to argue that lucky satisfaction of NDC is incompatible with knowledge. With this goal in mind, it could be tempting to think that there is a short and unproblematic step to this conclusion. The tempting line of thought is that the guard's actual belief (and not only her counterfactual belief) is epistemically irresponsible. If, as we are supposing, irresponsible belief is incompatible with knowledge, then the guard's *actual* belief would not amount to knowledge.

This line could be supported by an argument to the effect that the irresponsibility in the counterfactual situation somehow spreads to the actual one. It might be pointed out that the guard's irresponsible belief in the counterfactual situation is the result of a defective feature in her intellectual character, namely her (epistemically unjustified) deep contempt towards academics. It is this epistemically vicious aspect of her character that ultimately explains the irresponsibility of the counterfactual belief. But it might be argued that a subject with a defective or vicious epistemic character like this should be unable to form epistemically responsible beliefs – after all, the subject does have that vicious dispositional feature even in circumstances in which it is not actualized. According to this line, the guard's actual belief is irresponsible because it is the product of a vicious intellectual character.

But it is unclear that this tempting line is sound. If the guard's actual belief is irresponsible because it is the product of her vicious character, then nothing would prevent us from claiming that *all* of the guard's beliefs are irresponsible, for they all are the product of her intellectual character. But plainly it seems highly intuitive that the guard must be able to form *some* responsible beliefs even if *some* of her beliefs are irresponsible on the basis of the negative contribution by her contempt for scientists. In particular, it seems quite plausible to think that many of her perceptual beliefs could be regarded as epistemically responsible – for nothing in the case prevent us from construing the guard as an attentive and reliable (perceptual) believer, who assesses her perceptual evidence carefully before making a perceptual judgement. The defective aspect of her intellectual character is highly localized, and it is conceivable that it might play no role in the formation of many of her perceptual

beliefs. Why would a disposition that fails to be actualized in a given circumstance render the corresponding belief irresponsible? If the belief is the result of only virtuous aspects of the subject's character, then it is difficult to see why the belief should be regarded as irresponsible. If this reply is convincing, then it is possible to hold that the guard's belief in the actual situation is responsible.

It is unclear, then, that beliefs meeting NDC luckily will be epistemically irresponsible. Surely, my response in the previous paragraph need not be the last word in this dispute – a defender of the tempting line could very well develop a response to my objection. But instead of pursuing this line further, I set it aside for dialectical reasons. For I do not need to establish that beliefs that meet NDC luckily are epistemically irresponsible. In my view, these beliefs cannot constitute knowledge not because they are irresponsible, but because the kind of luck they exhibit is itself incompatible with knowledge.¹⁶ Surely, there might be cases where a belief that meets NDC luckily is, in addition, irresponsibly held; in which case its failure to count as knowledge will be overdetermined. But I want to argue that even if we think that beliefs that meet NDC luckily can be responsibly held, they still fail to constitute knowledge for the more fundamental reason that the kind of luck they exhibit is incompatible with knowledge. To this issue I turn now.

3.2 The Risk of Being Irresponsible

Now, the transition from the claim that a belief meets NDC luckily to the claim that it fails to constitute knowledge is far from being straightforward. For the transition requires showing that the kind of luck involved in luckily meeting NDC (hereafter I will refer to this kind of luck as “undefeated-doxastic luck”) is incompatible with knowledge. But one of the lasting lessons of discussions on luck after Gettier (1963) is that not all kinds of epistemic luck are incompatible with knowledge. This means that we cannot simply make an inference from the mere fact that a belief luckily satisfies a necessary condition for knowledge to the claim that this belief does not constitute knowledge. For instance, there is wide agreement in the claim that *evidential epistemic luck* is perfectly compatible with knowledge.¹⁷ This is the kind of luck involved in finding good evidence for *p* in a lucky fashion. Think, for instance, about a situation where I luckily overhear a conversation while waiting in line to

¹⁶ It is important to note that this does not mean that I defend an account where epistemic irresponsibility is the result of objective or external features. This would go against my claim that my view is consistent with Pritchard's (2005) account of epistemic irresponsibility, sketched in ft. 8 above. Lucky satisfaction of NCD is the result of (modal) objective or external features of the situation. Nevertheless, this does not render these (actual) beliefs epistemically irresponsible. Instead, the counterfactual beliefs are epistemically irresponsible. But the irresponsibility of those beliefs is readily explained by subjective or internal factors that obtain in the counterfactual situation. The guard's belief in the counterfactual situation is irresponsible because the guard would be disregarding cogent evidence that she possesses. This means that, from her epistemic perspective, her belief would be true by luck – this is the “reflective epistemic luck” identified by Pritchard (2005).

¹⁷ See Unger (1968) for an early discussion of this point, as well as Pritchard (2005) for a more detailed exploration of the different kinds of epistemic luck and their impact with respect to knowledge possession.

order coffee, which gives me evidence that I would otherwise lack for thinking that the Central Bank will announce an interest rate cut today. Too easily I could have failed to overhear the conversation, so my acquiring the relevant evidence is a lucky event, as it is my forming the corresponding belief. But it seems intuitive that this is compatible with my coming to know that the Central Bank will announce an interest rate cut. Luckily finding evidence seems compatible with knowing the propositions supported by that evidence.

But if not all kinds of luck involved in (true) belief acquisition are incompatible with knowledge, we need to say more in support of the idea that undefeated-doxastic luck is incompatible with knowledge. This is an important task, for defenders of NDC could take refuge in the claim that meeting NDC luckily does threaten the corresponding beliefs' status as knowledge. After all, they could argue, it is *actual* failure to satisfy NDC that is incompatible with knowledge. A potential explanation as to why some epistemologists could find this line of reasoning attractive is that they might think that any alternative is bound to be overly demanding, and would place an unrealistic burden on most ordinary knowers. They might worry that any attempt to eliminate the kind of luck identified above would require subjects to be *constantly* and *actively* monitoring their situation for the presence of potential defeaters. I will address this concern in detail in §4. For the time being I just want to point out that my proposal will not impose those unrealistic demands. I shall argue that subjects can exhibit an adequate responsiveness to defeaters without constantly and actively monitoring their situation.

My claim is that undefeated-doxastic luck is problematic in a way similar to that of the *veritic luck* distinctive of Gettier cases. I shall argue that their structural similarities give us a principled reason to think that both are incompatible with knowledge. A true belief is veritically lucky when it could easily have been false, i.e., when there are nearby counterfactual situations where the same belief formed in a similar fashion turns out to be false. Take the classic example of seeing a hairy dog afar in a field and forming on this basis the belief that there is a sheep in the field. In the standard example the belief turns out to be true but only because a sheep (which could easily not have been there) is somewhere in the field hidden from view.¹⁸ This belief is *luckily true*. There is wide agreement that this belief (as well as other beliefs exhibiting the same kind of luck) does not constitute knowledge. This is why, for instance, safety-based responses to the Gettier problem focus on developing conditions for knowledge which ensure that the counterfactual situations where one's belief would be false are modally far away.¹⁹

We have examined two different kinds of epistemic luck – evidential and veritic – and we have seen that there is wide agreement in thinking that the former is compatible with knowledge, whereas the latter is not. Taking this differential diagnosis as a *datum* one might wonder what exactly explains it. I shall argue that an adequate explanation provides us with reason to think that undefeated-doxastic luck is also incompatible with knowledge.

¹⁸ This example is owed to Chisholm (1966).

¹⁹ See Sosa (2009), Williamson (2000), Pritchard (2005).

Let me begin my case by reflecting on the widely shared idea that a general aim of our cognitive lives should be that of attaining true beliefs *and* avoiding false ones.²⁰ Perhaps this maxim gets its support from a more general one, to the effect that cognitive agents should aim to attain epistemic goods *and* avoid epistemic failures. This more general maxim would allow us to defend specific maxims such as “aim to attain knowledge and avoid beliefs that fall short of being knowledge” or “aim to form responsible beliefs and avoid forming irresponsible ones”. The importance of recognizing both the positive goal (attain epistemic goods) as well as the negative one (avoid epistemic failures) has a venerable history, for it is clearly recognized by Descartes, as he repeatedly points out in the *Meditations* that even if we turn out to be unable to attain clear and distinct conceptions of many of the issues that we took ourselves to know before embarking on the self-reflective enterprise, we still have a path available to save our integrity as cognitive agents – for we can always avoid failure by suspending judgement.²¹

Arguably, this maxim which is meant to provide a very general guide for our cognitive lives can explain our datum – i.e., our differing intuitions regarding veritic and evidential luck. Shortly, the explanation is that evidential luck is compatible with knowledge because the subjects whose true beliefs are evidentially lucky are in *no risk* of forming a *new* false belief; whereas subjects whose true beliefs are veritically lucky *incur* in such risk. I assume a conception of risk which construes it as a sister notion to that of luck.²² That is, there is a risk that a state obtains or that an event occurs when, despite failing to occur or obtain in the actual circumstances, there are modally nearby counterfactual circumstances where the event occurs, or the state obtains. The idea, then, is that in cases of veritic luck subjects runs the risk of incurring in an epistemic failure (false belief), whereas in cases of evidential luck the subject does not run such risk. But why is this risk incompatible with knowledge? A natural proposal is that the specific risk incurred by veritic luck (namely, false belief) is itself incompatible with knowledge. That is, veritic luck involves the risk of false belief, which amounts to the risk of failing to know (for a *fortiori* a false belief cannot constitute knowledge). In contrast, this risk is absent in cases of evidential luck. My proposal is that this difference explains why we have the intuition that evidential luck is, whereas veritic luck is not, compatible with knowledge.

Let me illustrate these claims with some of the examples previously discussed. The modally nearby situations where I fail to overhear the conversation about the Central Bank are not situations where my failing to overhear it (by itself) leads to the formation of a new false belief (about the interest rate).²³ In other words, in the counterfactual situation where I fail to overhear the conversation, I am not epistemically worse-off than I was in the actual world in the moment previous to my overhearing

²⁰ For the early standard formulation of this idea see James (1956).

²¹ See, e.g., Descartes (1996), p. 43.

²² For the notion epistemic risk and how it is closely connected to the notion of luck see Pritchard (2016).

²³ Of course, the case can be further complicated to make it such that failing to overhear that conversation does lead me to a false belief. I treat this kind of case in detail below.

the conversation. For even if I already had a false belief on the matter (e.g., that the Bank will not lower interest rates), my failing to overhear the conversation does not make my situation worse than it was before (I simply continue having a false belief I previously had anyway). This means that in this case there is no risk of being *epistemically worse-off*, for there is no risk of forming a *new* false belief on the matter. In contrast, in the sheep example the modally nearby situations where the sheep is not in the field (and I am deceived by the hairy dog) are situations where I do end up with a new false belief (i.e., that there is a sheep in the field) that I did not previously have. This means that the counterfactual situation where no sheep is present is a situation where I am epistemically worse-off than I was before encountering the problematic evidence. In both cases there is a sense in which I am lucky that I have a true belief, but only in one of them it is also the case that I formed the belief in a manner that risked my ending up with a new false belief.

If this diagnosis is correct, then, with minor unproblematic further assumptions, the diagnosis can be extended to cover cases of undefeated-doxastic luck. I assume that the general epistemic maxim supports the following specification: cognitive agents should endeavour to form and sustain responsible beliefs, and avoid forming and sustaining irresponsible ones. I maintain that undefeated-doxastic luck involves running the risk of believing irresponsibly. In this respect, this kind of luck is more similar to veritic luck than to evidential luck. As my central case has been described, the security guard's situation is such that she runs the risk of sustaining an irresponsible belief, for there is a modally nearby situation where the guard's belief is irresponsible (by virtue of failing to meet NDC). This means that the relevant counterfactual situation is one where the guard is epistemically worse-off than she is in the actual world. Thus, just like veritic luck, undefeated-doxastic luck involves running the risk of forming a belief in a manner which would give rise to an *epistemic failure*. Moreover, in both cases the risk involves an epistemic failure that is *inconsistent* with knowledge – for just as a belief's falsity is inconsistent with knowledge, so is irresponsibility. This would provide us with a principled *prima facie* reason to think that beliefs that meet NDC luckily cannot constitute knowledge – the reason being that undefeated-doxastic luck is structurally identical to veritic luck.

3.3 Objections and Replies

Before articulating my own positive proposal about how to address this pernicious form of epistemic luck let me address two potential objections to the diagnosis I have just put forward. First, recall that I left open the possibility of maintaining that the security guard's belief is *irresponsible* in the actual situation (§3.1). This could be exploited to try to raise a challenge for my proposal. It could be argued that the counterfactual situation is one where the guard is not epistemically worse-off, for she is already epistemically irresponsible in the actual situation. Therefore, she would not run the risk I identified above. There are two things to say in response to this objection. First, if we assume that the guard's actual belief is irresponsible then all the better for my position, for this gives us a shortcut to my desired conclusion, namely that a belief meeting NDC luckily does not constitute knowledge. Second, if

you think that the guard's belief is irresponsible in the actual scenario it only follows that the guard is not epistemically worse-off with respect to doxastic responsibility. Nevertheless, there is another respect in which she clearly is worse-off: by virtue of meeting NDC luckily she runs the risk of failing to meet NDC – and failing to meet NDC is also an epistemic failure that is incompatible with knowledge.

Second, it might be argued that I have cherry-picked the examples I employ to illustrate evidential luck in order to favour the diagnosis that I advance, and that I have left out cases which raise difficulties for my diagnosis. In particular, it might be argued that there are cases of evidential luck where subjects *do* run the risk of having a new false belief. This would render at least some cases of evidential luck more similar to cases of veritic luck than I take them to be. Consider barn-façade-style cases.²⁴ There is a sense in which subjects are lucky to have the good evidence they in fact have for thinking that a barn is before them – for very easily they could be facing a fake barn and lack the good evidence that they have. But in these cases, subjects seem to run the risk of forming a new false belief (i.e., to think that there is a barn when they face a façade). Accordingly, these cases give rise to mixed intuitions – some epistemologists favour the idea that the luck involved is inconsistent with knowledge, while others maintain that this kind of luck is perfectly compatible with knowledge. Thus, this objection goes, we have cases of evidential luck where subjects seem to run the risk of forming a new false belief *and* where at least some philosophers have the intuition that this kind of luck is incompatible with knowledge. Both claims go against my favoured diagnosis.

My response to this worry is that these cases do not straightforwardly exhibit *only* evidential luck; in fact, they seem to involve a mixture of evidential and veritic luck. This could explain why these cases involve the risk of forming a new false belief, which in turn may partly explain the mixed intuitions that they elicit. The very concept of evidential luck involves nothing more than the evidence *e* being acquired in a lucky manner. For *e* to be acquired luckily it is only needed that in a nearby counterfactual situation the subject *does not* acquire *e*. This much suffices to make the case one of evidential luck. Clearly, failure to acquire *e* is different and does not entail the acquisition of alternative (and perhaps misleading) evidence *e'*. The acquisition of alternative evidence in a nearby counterfactual situation is an additional part of the story, unnecessary for the situation to count as a case of evidential luck. Accordingly, the overheard conversation case sketched above is a case where no alternative misleading evidence is in the offing, which makes the case one of pure evidential luck. But barn-façade-style cases *do* involve the acquisition of misleading evidence in nearby counterfactual situations. As a result, barn-façade cases involve something more than just evidential luck. As a matter of fact, they seem to involve a type of veritic luck, for there is a clear sense in which it is partly a matter of luck that the actual belief is true – very easily it could have been false. Accordingly, the relevant counterfactual situation is one where the subject forms a new *false* belief. Thus, the subject in barn-façade County is lucky in finding herself in front of the only barn in the County (evidential luck), but she is also lucky in that her belief turns out to be

²⁴ The original barn-façade case can be found in Goldman (1976), although he attributes it to Ginet.

true (veritic luck).²⁵ This means that these cases, being "impure" cases of evidential luck do not show that my original diagnosis is incorrect. The risk of forming a new false belief in barn-façade cases can be naturally attributed to veritic luck. In turn, this might explain the mixed intuitions that these cases elicit. In particular, the presence of veritic luck might explain the intuition that the luck involved in these cases is inconsistent with knowledge.

4 A Solution: Being Responsive to Defeaters

Undeclared-doxastic luck ultimately is the result of an inadequate treatment of the threat posed by epistemic defeaters. NDC is too weak to address this threat because it can be met fortuitously, through no effort from the believer. When the condition is met fortuitously, the subject is entirely vulnerable to the presence of defeaters in nearby counterfactual situations, an epistemic risk that cannot be tolerated by knowledge.

As a result, the main task in what follows should be to develop a condition which adequately keeps the threat of defeaters at bay, and which is informed by the foregoing discussion regarding undeclared-doxastic luck. A lesson from the security guard example is that she falls prey to this kind of luck because she is utterly *unprepared* to respond adequately to certain defeaters. A straightforward way to fix this problem is to make sure that knowers are not unprepared in this way. Thus, I propose a more modally robust condition according to which subjects ought to be appropriately responsive to epistemic defeaters. As we shall see, this condition cannot be met fortuitously, since meeting it involves being *disposed* to react appropriately if defeaters arise. It follows that subjects should be able to address defeaters in counterfactual situations even if none are present in the actual situation.

As a first pass I advance the following requirement in substitution of NDC:

Defeaters-responsiveness: a belief that *p* constitutes knowledge only if the believer is adequately responsive (i.e., has a disposition to respond appropriately) to epistemic defeaters against *p*.

It should be clear that the notion of "being adequately responsive to defeaters" is the one doing the heavy lifting in this condition. It is imperative, then, to flesh it out in detail.

On the proposal I have in mind, responsiveness to defeaters takes two different forms depending on the kind of defeater at stake. On the one hand, in order to be

²⁵ It is precisely on the basis of this type of considerations that Pritchard (2015) construes the kind of luck present in barn-façade cases as a form of veritic luck. He labels it "veritic *environmental* luck" – for this is a situation where the belief is true by luck because of the specific layout of the environment the subject is in. It differs from the kind of veritic luck that obtains, e.g., in the sheep example, labelled "veritic *intervening* luck". The latter is a kind of luck where something "intervenes" between the fact and the relevant belief; in contrast, in veritic environmental luck nothing "intervenes" between fact and belief. For the reasons sketched above, I favour a construal where this barn-façade cases involves a mixture of evidential and veritic luck.

appropriately responsive to *psychological* defeaters subjects should be able to either neutralize them *or* to suspend judgement on the threatened proposition. On the other hand, in order to be appropriately responsive to normative defeaters subjects should (in addition to the previous condition) have strategies in place to make sure that they would reliably identify normative defeaters if they were to arise.

In a nutshell, this is what the requirement amounts to. Recall that earlier I pointed out that some epistemologists might worry that the requirement I propose here moves dangerously towards hyper-intellectualism – that is, the worry that a theoretical account of knowledge imposes unrealistic, overly demanding, conditions for subjects to attain knowledge. For instance, they might worry that for one's belief to be immune to undefeated-doxastic luck subjects would need to constantly and actively monitor their epistemic landscape looking for potential defeaters. If this was what it takes to avoid undefeated-doxastic luck, then I would agree that we could hardly have any knowledge, for I concede that most subjects would be unable to meet the condition most of the time. My initial exposition of the defeater-responsiveness condition was rather meagre. Let me now spell out the two forms of the requirement in more detail while addressing the worry of hyper-intellectualism.

On the proposal I have set forth, subjects can be prepared to respond adequately to the presence of *psychological defeaters* without having to do much. Subjects can address (psychological) defeaters adequately by retorting to suspension of judgement. I contend that this is not something that is overly demanding for most subjects most of the time – for it merely amounts to being disposed to suspend judgement when one recognizes that there is a defeater for one's belief (and one lacks a good response to it). In contrast, our security guard, recall, did not even have this disposition. That this is always an available way to be responsive to defeaters shows that this form of the requirement is not overly demanding. But, in my view, following Descartes (1996), suspension of judgement is merely a fallback option: subjects would be better-off epistemically if they were able to neutralize potential defeaters for their beliefs. This is another, better, way in which subjects can exhibit defeater-responsiveness. Some agents might be better prepared to neutralize potential defeaters than others, and most likely this will result in a better epistemic position overall. For instance, an experienced ornithologist might be able to neutralize many potential defeaters that could arise for her belief that there is a goldfinch before her. But, in contrast, I might be unable to neutralize most potential defeaters for my belief that the bird is a goldfinch – and, thus, in most cases I would have to retort to suspension of judgement. My contention is that in both cases the defeater-responsiveness condition would be met because both subjects are adequately responsive to psychological defeaters – although obviously the expert ornithologist would be better positioned to gain knowledge than I would be.

But in the case of normative defeaters a similar line of response against the hyper-intellectualism worry is unavailable. For in these cases there is not a homogeneous fallback option which requires very little from cognitive agents (like suspension of judgement does). In this case, the condition requires from subjects to have strategies in place to avoid being unaware of information that they should be aware of. Now, it is impossible to specify what form these strategies should take in abstraction of the operative circumstances. The epistemic context does play an important role in

specifying the condition for each subject. Importantly, different subjects (by virtue of their different situations) can differ in what kind of information they have an obligation to be aware of. A person working as an investment manager needs to be far more up to date with the economic news than I need to be. As a result, the strategies the investment manager needs to deploy to be responsive to normative defeaters on this topic would take a very different form than mine. Given her epistemic obligations, she has to make sure that she knows the economic news as soon as possible; whereas I, given my obligations, need not do any of the like. It certainly seems that my defeaters-responsiveness condition would translate into some fairly demanding requirements for the investment manager. I think this is exactly the right result and that it does not constitute a form of hyper-intellectualism for two reasons. First, there are independent reasons to think that subjects like the investment manager have rather steep epistemic obligations to meet. Those reasons in fact motivate the very idea of normative defeaters – information that subjects *should* be aware of. I contend that responsiveness to normative defeaters does not add any additional requirement which is not already accepted by virtue of accepting the very notion of normative defeaters. If you think that the investment manager should be aware of the daily economic news, then you must think that she needs to implement reliable strategies to satisfy this obligation – on pain of epistemic irresponsibility. Second, not everyone is an investment manager and not everyone has the same epistemic obligations. Which means that not everyone is required to implement fairly demanding strategies to monitor certain aspects of their epistemic landscape. Some epistemic agents have rather limited obligations, which means that they need not constantly and actively monitor their landscape for defeaters. But others do, and for them a requirement to monitor their landscape is not an overly demanding requirement – for they already incur in that obligation for independent reasons.

Thus, defeaters-responsiveness is not an overly demanding requirement for knowledge. When it comes to psychological defeaters subjects can always retort to suspension of judgement – a not overly-demanding requirement. And when it comes to normative defeaters, not all subjects are required to constantly and actively monitor their landscape for defeaters since they might lack the obligation to implement such a strategy. When subjects are required to implement such strategies, this is grounded on independent reasons about their specific epistemic obligations.

However, even if it is accepted that this is not an overly demanding requirement, it could be argued that it leads to a seemingly counter-intuitive consequence, namely that the investment manager has to work harder than the non-expert to get the same piece of knowledge. In the face of it, this seems counter-intuitive. But once we pay attention to further consequences of the *epistemic work* carried out by the investment manager, we can see that this is not so counter-intuitive after all. Consider the proposition (*p*): “the interest rates are currently low (i.e., 0.25%)”. On my view, a non-expert might get to know that *p* on the basis of weeks-old news reports that she remembers. Since she is not obliged to be up to date she can rely on this kind of information. In contrast, the investment manager can only know that *p* if she successfully keeps tracks of these announcements and is aware of the most recent one. But notice that the extra epistemic effort incurred by the investment manager does come with some epistemic rewards: her knowledge that *p* is more secure than the

non-expert's. It is based on a method that is more sensitive to incoming defeaters, so she is able to address them more efficiently and change her doxastic attitudes accordingly. In other words, the expert's knowledge enjoys some epistemic protection that the non-expert's knowledge lacks (this view mirrors some suggestions made by Plato in *Meno* (2009)). Thus, the expert does work harder to get the same piece of knowledge, but the extra work is explained by the improved protection acquired by the expert.²⁶

Before wrapping up, I would like to consider one final objection to the view that I have advanced. It could be argued that my argumentative strategy for the claim that undefeated-doxastic luck is incompatible with knowledge could be re-deployed with respect to defeater-responsiveness and engender a second-order anti-luck requirement. In turn, this could arguably lead to global scepticism. The idea goes as follows: it could be a matter of luck that it is not a matter of luck that my belief that p is responsibly held. But if this kind of second-order luck obtains, then my belief that p would not be knowledgeable, for it would be a matter of luck that it satisfies a necessary condition for knowledge. This would mean that this belief incurs in an inadmissible type of epistemic risk. To fix this problem, we would require a second-order anti-luck requirement, which would require the belief that p to be responsibly held in an even wider class of counterfactual situations. However, since this procedure could be iterated indefinitely, the reasoning goes, we would end up with a requirement to the effect that the belief ought to be responsibly held in *all* counterfactual circumstances.²⁷ But that is clearly an overly demanding requirement, which many beliefs would fail to meet. Thus, the threat of scepticism arises.

This is indeed a powerful and serious challenge which ought to be adequately addressed if the view that I defend here is to be plausible at all. However, the task of developing a detailed and cogent response to this challenge must be left as an outstanding task, for here I lack the space to develop such a response in detail. However, I would like to outline the beginning of a potential response to this challenge. We must begin by noting that the challenge above depends on the idea that moving from a first-order anti-luck requirement to a higher-order anti-luck requirement has the effect of widening the class of counterfactual situations that are relevant for the corresponding epistemic evaluation. After a number of "jumps" into higher-order requirements, we reach a maximally wide set of counterfactual situations, which in turn engenders the sceptical threat. The core of my response depends on rejecting the assumption that a "jump" to a higher-order requirement widens the class of counter-factual situations relevant for epistemic evaluation. On my view, the class

²⁶ An interesting consequence of this view is that there are important differences in the epistemic obligations of subjects, depending on the social or epistemic position that they occupy. This is a relativity phenomenon within epistemology that has not been discussed in much detail so far. Importantly, it differs from the relativity that appears in standard pragmatic encroachment and contextualist positions. I would like to thank an anonymous reviewer of this journal for calling to my attention the seemingly counter-intuitive consequence cited in the body of the paper as well as the interesting consequence mentioned in this footnote.

²⁷ A further worry is that a similar reasoning could be developed for every necessary condition on knowledge, such as truth, warrant or safety.

of counterfactual situations if fixed by the target belief under evaluation. This class must be limited to the modal vicinity of that target belief because we are interested in assessing the epistemic risks of *that* belief. And to do that we only need to look into the counterfactual situations nearby *that* belief.

The widening of the class of counter-factual situations described in the objection would occur if we were moving from assessing a first-order belief (e.g., that *p*) to assessing a second-order belief (e.g., that the belief that *p* is responsibly held). But the type of higher-order anti-luck requirement engendered by my arguments does not produce a shift from first- to second-order belief. Since my view does not entail such a shift, my view does not produce the widening of the class of counter-factual situations that seems to lead to the sceptical consequences identified above. Recently, some have argued that modal accounts of safety fall prey to this type of sceptical challenge (Greco, 2016) precisely on the premise that moving from first-order safety to higher-order safety involves a jump to higher-order beliefs. If my contention (which needs to be fleshed out) is correct, the required premise is not available to mount the sceptical challenge. Now, this is merely the sketch of a potential response, which ought to be fleshed out elsewhere. However, it is important to note that my response reaches similar results to those advanced by Greco himself (2016). On his proposal, we ought to move away from a modal closeness account of safety and favour a context-sensitive notion of robustness. Nevertheless, if the proposal sketched here is successful, we would be able to achieve similar results without abandoning a modal closeness account of safety.²⁸ As mentioned, this is merely an outline of a response to the sceptical challenge. For the view defended here to be plausible, we need to develop this response in detail. Due to space constraints, I must leave that task for a different occasion.

5 Why Does it All Matter? Some Consequences of the View

If the foregoing argument for Defeater-responsiveness is successful, this has important consequences for our understanding of knowledge, for we are provided with an improved version of a necessary condition for propositional knowledge. It has been shown that the no-defeaters condition is insufficient to block the threat that epistemic defeaters bring to the table. To address this problem, I have argued that, additionally, we need to require subjects to be ready to react adequately if defeaters were to arise.

But, as mentioned in the introduction, the proposal I have set forth here matters and has implications beyond providing us with an improved formulation of a necessary condition for propositional knowledge. It tells us something important about the epistemic agency that subjects need to exercise if they are to qualify as knowers. On certain "narrow" conceptions of epistemic agency, this is limited to the assessment of one's evidence and the adjustment of one's doxastic states in response to those assessments

²⁸ I would like to thank an anonymous reviewer of this journal for alerting me about the existence and importance of this sceptical threat.

(e.g., Christensen, 2007; Matheson, 2016; Sosa, 1985, 2015). In contrast, I suggest that the proposal advanced in this paper provides us with good reasons to defend a wider conception of epistemic agency.

On the narrow conception of epistemic agency, when conditions are auspicious subjects need not do much in order to gain knowledge – i.e., all they have to do is assess their evidence and form the corresponding judgement (Sosa, 2015: 45). Certainly, in less auspicious circumstances, subjects might be required to do much more – e.g., they might need to put some “effort and persistence” (Sosa, *idem*) in order to obtain hard to get evidence. Nevertheless, for Sosa, this *extra* effort and persistence is *dispensable*, for subjects can gain lots of knowledge *without* going through the extra effort. As such, the agency involved in the extra effort is not strictly *epistemic* agency – for knowledge can be gained without its involvement.

Against this narrow view, it could be argued that the agency involved in the “extra effort” *is* epistemic – but I will not pursue that line here. Instead, let us concede that only agency that is *indispensable* for knowledge acquisition counts as genuine epistemic agency. Even with this strict criterion for epistemic agency in place, the proposal I have advanced provides us with reasons to defend a *wider* conception of epistemic agency. For, as we have seen, defeater-responsiveness *is* necessary for propositional knowledge. In turn, defeater-responsiveness seems to involve agency (therefore, epistemic agency) in an important way. For the dispositions that subjects ought to have in order to satisfy defeater-responsiveness are *dispositions to act*. These include dispositions to suspend judgement, neutralize defeaters and monitor one’s epistemic landscape. Even if in a given occasion (i.e., in the most auspicious of circumstances) these dispositions are not actualized, they remain *agential* dispositions, and both their possession as well as their exercise should be regarded as genuine parts of our epistemic agency, since they are necessary for propositional knowledge.

To conclude, let me review what has been achieved. I have argued in detail that a no-defeaters condition is insufficient to address the threats that epistemic defeaters bring about. In particular, it is unable to block undefeated-doxastic luck – a kind of luck that, I argue, is incompatible with knowledge. In response, I proposed incorporating a defeater-responsiveness condition for knowledge. I briefly unpacked that condition and argued that it does not impose overly demanding requirements for knowledge. The main result of this discussion is the development of an improved formulation of a necessary condition for knowledge. In turn, this has consequences for discussions on the extent of human epistemic agency: I briefly sketched an argument against a narrow conception of epistemic agency. On my view, epistemic agency goes beyond merely assessing evidence and making a corresponding judgement. The dispositions involved in meeting the defeater-responsiveness requirement should be considered part of our epistemic agency, for they are necessary for propositional knowledge acquisition.

Acknowledgements I would like to thank Martín Abreu Zavaleta, Diego Rodríguez, Oscar Piedrahita, Alejandro Vesga, Shawn Wang, and two anonymous referees of this journal for very helpful discussions and suggestions on earlier versions of this paper. Special thanks to Florencia Rimoldi and Miguel Ángel Fernández for very useful and detailed feedback on an even earlier version of this paper. Thanks to Angeles Erana for her supervision during my postdoctoral stay at UNAM. This work was made possible thanks to the financial support of the postdoctoral program of the National Autonomous University of Mexico at the Institute for Philosophical Research. In addition, this publication was made possible

through the support of the grant #61255 from the John Templeton Foundation. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the John Templeton Foundation.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anaya, A. (2021). Perceptual presentation and the myth of the given. *Synthese*. <https://doi.org/10.1007/s11229-021-03124-5>
- Bergmann, M. (2000). Deontology and defeat. *Philosophy and Phenomenological Research*, 60(1), 87–102.
- Bergmann, M. (2005). Defeaters and higher-level requirements. *The Philosophical Quarterly*, 55(220), 419–436.
- BonJour, L. (1980). Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5(1), 53–73.
- Chisholm, R. (1966). *Theory of Knowledge*. Prentice-halls, .
- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116(2), 187–217.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- de Grefte, J. (2020). Towards a hybrid account of luck. *Pacific Philosophical Quarterly*, 101, 240–255.
- Descartes, R. (1996). *Meditations on First Philosophy. With Selections from the Objections and Replies*, Ed. And translated by John Cottingham. Cambridge: Cambridge University Press.
- Engel, P. (2009). Epistemic responsibility without epistemic agency. *Philosophical Explorations*, 12(2), 205–219.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123.
- Goldberg, S. (2018). *To the best of our knowledge*. Oxford University Press.
- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73, 771–791.
- Goldman, A. (1979). "What is Justified Belief," in George Pappas (ed.), *Justification and Knowledge*. D. Reidel, 1–23.
- Greco, D. (2016). Safety, explanation, iteration. *Philosophical Issues*, 26, 187–208.
- James, William. 1956. *The Will to Believe and Other Essays in Popular Philosophy*. Dover Publications.
- Kornblith, H. (1983). Justified Belief and Epistemically Responsible Action. *The Philosophical Review*, 92(1), 33–48.
- Lackey, J. (2005). Memory as a generative epistemic source. *Philosophy and Phenomenological Research*, 70(3), 636–658.
- Lasonen-Aarnio, M. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24(1), 1–21.
- Matheson, J. (2016). "Disagreement and the Ethics of Belief," in James Collier (ed.), *The Future of Social Epistemology: a Collective Vision*. Lanham, Maryland: Rowman and Littlefield, 139–148.
- McDowell, John. 1993. "Knowledge by Hearsay. In B. K. Matilal and A. Chakrabarti (eds.), *Knowledge from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*. Dordrecht: Kluwer Academic Publishers. Reprinted in John McDowell, 1998, *Meaning, Knowledge, and Reality*. Cambridge Mass.; London: Harvard University Press, 414–444.
- McDowell, J. (1996). *Mind and World*. Cambridge Mass.; London: Harvard University Press.
- McDowell, J. (2013). Tyler Burge on Disjunctivism (II). *Philosophical Explorations*, 16(3), 259–279.
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

- Plato. . (2009). *Meno in Meno and other Dialogues* translated by Robin Waterfield. Oxford University Press.
- Pollock, J., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Rowman and Littlefield.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2015). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172(1), 93–111. <https://doi.org/10.1007/s11098-014-0374-0>
- Pritchard, D. (2016). Epistemic risk. *The Journal of Philosophy*, 113(11), 550–571.
- Setiya, K. (2013). Epistemic agency: some doubts. *Philosophical Issues*, 23(1), 179–198.
- Smithies, D. *Forthcoming*. The Epistemic Function of Higher-order Evidence. In Paul Silva & Luis Oliveira (eds.), *New Essays on Propositional and Doxastic Justification*. Routledge.
- Sosa, E. (1985). Knowledge and intellectual virtue. *The Monist*, 68(2), 226–245.
- Sosa, E. (2009). *A virtue epistemology: Apt belief and reflective knowledge*, (Vol. I). Oxford University Press.
- Sosa, E. (2010) *Knowing Full Well*. Princeton University Press.
- Sosa, E. (2015). *Judgement and agency*. Oxford University Press.
- Unger, P. (1968). An analysis of factual knowledge. *The Journal of Philosophy*, 65(6), 157–170.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.