



# Epistemically Vicious Knowledge

Spencer Paulson<sup>1</sup>

Received: 15 August 2023 / Accepted: 2 August 2024  
© The Author(s), under exclusive licence to Springer Nature B.V. 2024

## Abstract

I will present a novel argument that there can be epistemically vicious knowledge. In the kind of case that interests me, the subject knows not despite but rather because of her vice. It is generally agreed that some kinds of epistemic luck doesn't undermine knowledge. For instance, being lucky not to have misleading evidence doesn't undermine knowledge. I will argue that this doesn't change when the avoidance of misleading evidence depends on the subject's vice. It does not prevent her belief from being both safe and sensitive, nor does it have any bearing on whether her justifying reasons are ultimately defeated. To illustrate my point, I will draw a parallel with a certain way of thinking about resultant moral luck. I will then go on to consider the ramifications for virtue-theoretic analyses of knowledge and the doctrine of normative defeat.

**Keywords** Epistemic Luck · Epistemic Vice · Virtue Epistemology · Moral Luck

Here is an uninteresting case of epistemically vicious knowledge: a subject misused a telescope because of intellectual laziness but her misuse caused her to acquire knowledge that the lens is out of focus. Can there be interesting cases? Alfano and Levy (2019) argue that there can be at the social level: cumulative cultural evolution, which requires the transmission of knowledge, depends on the epistemic vices of individuals. Baehr (2014a, b) argues that there can be at the individual level. More specifically, he argues that one can have knowledge that *p* *despite* viciously inquiring into whether *p*.

---

✉ Spencer Paulson  
spencerpaulson2023@u.northwestern.edu

<sup>1</sup> Northwestern University, Kresge Hall 3512, 1880 Campus Drive, Evanston, IL 60208, USA

In this paper I will argue, like Baehr, for the possibility of epistemically knowledge at the individual level.<sup>1</sup> Unlike Baehr, I will argue that a subject can come to know that *p* not despite but rather *because* her inquiry into whether *p* was epistemically vicious. If successful, my argument will present an even deeper worry about the prospects for virtue-theoretic analyses of knowledge than the one that arises from the prospect of knowledge despite vice. In the cases of interest to Baehr, the fact that you truly believe *p* doesn't depend on the vicious inquiry. In the cases that interest me, the subject truly believing *p* depends on the vicious inquiry (i.e., it depends on at least one part of that inquiry being vicious, not necessarily all or most of it).<sup>2</sup> So, the virtue theorist can't dismiss the vice as explanatorily superfluous.

Virtue theorists claim that a subject's true belief is creditable to them just in case they believed truly because they formed their belief virtuously.<sup>3</sup> Different virtue theorists give different accounts of virtue and different accounts of the "...because of..." relation, but I will argue that the differences do little to mitigate the problem. The reason is that it is unclear why the subject deserves credit if their true belief is no less dependent on their inquiring viciously than it is on their inquiring virtuously. Even if their believing truly depends on virtue and vice in different ways, it is unclear why the subject deserves credit when both kinds of dependence are explanatorily significant. Since the credit account holds that knowledge is true belief for which the subject deserves credit, it is consequently unclear why only one matters for knowledge.

My argument goes as follows. I will assume you can be lucky to have the evidence you do and still have knowledge. Knowledge-undermining luck involves the relation between the evidence (or a belief based on it) and the truth, not whether you have the evidence in the first place. This is relatively uncontroversial. I will argue that it is no less plausible to think you can be lucky not to have misleading evidence without this luck undermining your knowledge. I will then argue that even if the explanation of why you were lucky not to have misleading evidence is that you were epistemically vicious, you might still have knowledge. That is the main argument. To illustrate the general point that one can be lucky to have a status and, even if one is lucky because of vice, maintain that status, I will develop the following parallel case. I assume (purely for the sake of exposition and not as a premise in my reasoning<sup>4</sup>) that there is such a thing as resultant moral luck. I then argue that resultant moral luck is compatible with that luck largely being explained by the manifestation of an epistemic vice. I then show that the epistemic case is structurally identical to the moral case.

What this ultimately shows is that one's knowledge can be just as attributable to vice as virtue. I will argue in the final section of the paper that this is a problem for virtue theoretic analyses of knowledge according to which knowledge is true belief creditable to the subject because the fact that she believes truly is explained by her

<sup>1</sup> In fact, I will use an example somewhat similar to his vicious scientist (see Baehr, 2014a: 137) to make my point.

<sup>2</sup> Thanks to an anonymous referee for raising this point.

<sup>3</sup> Cf. Zagzebski (1996, 2003, 2008), Kvanvig (2003), Greco (2003, 2010, 2012), Sosa (2007) and Turri (2011).

<sup>4</sup> As will be shown, even if there isn't resultant moral luck in the controversial sense, one could still be lucky not to have produced a worse outcome than one in fact did. This is really all the parallel requires.

cognitive virtues (or something along these lines). In short, it may be that knowledge *despite* vice doesn't compromise this order of explanation because vice doesn't play the same explanatory role that virtue does. However, if one knows *because* of vice, then it is no longer why the credit-worthiness resulting from the explanatory role of virtues isn't canceled out by the equally important explanatory role played by vice.

In section (I), I discuss one kind of luck that doesn't undermine knowledge: being lucky to have evidence. I then go further and argue it is just as plausible to think you can be lucky not to have misleading evidence without this undermining your knowledge. Intuitive examples include already well-known cases in epistemology. In section (II) I turn my attention to resultant moral luck. I discuss it just enough to argue that resultant moral luck is compatible with moral vice figuring prominently in the explanation of why the subject got lucky. In section (III) I return to the epistemic case and argue that something analogous holds there. Just as you can have your moral luck explained by moral vice, your epistemic luck can be explained by epistemic vice. In section (IV) I discuss the ramifications for virtue epistemology.

## 1 Luck and Knowledge

It is generally agreed that Gettier subjects don't know because of knowledge undermining luck.<sup>5</sup> There is not widespread agreement about how to give an account of knowledge undermining luck. I will briefly discuss two accounts here: one that has emerged from the defeasible reasoning literature<sup>6</sup> and another that attempts to understand epistemic vocabulary in modal terms (i.e., safety and sensitivity). I won't commit myself to either one here, but I will try to show that they both have some promise for explaining why some kinds of epistemic luck undermine knowledge and others don't.

Let us first consider the defeasible reasoning account. Ordinary reasoning, unlike the kind of reasoning we use in deductive logic, is non-monotonic. In deductive logic, adding premises will never result in a decrease of consequences you can infer. In ordinary reasoning, this isn't the case. If I know that today is Wednesday, that allows me to infer (which is not to say prove) that the leasing office is open. But if I then learn that today is a holiday, I can no longer infer that. The fact that today is a Wednesday is a *prima facie* reason to believe that the leasing office is open. But it may not be an all-things-considered reason to believe that. For it to be an all-things-considered reason, it must not be defeated. Defeaters are *prima facie* reasons to give up a belief, they are the mirror images of *prima facie* reasons.<sup>7</sup> Since defeaters are *prima facie* rather than all-things-considered reasons to give up a belief, they too can be defeated. Sometimes defeater-defeaters restore your original justifying reasons. I might believe that the wall is red because of my perceptual experience, then receive

<sup>5</sup> The few who disagree include Weatherston (2003), Olsson (2015) and Weinberg (2017).

<sup>6</sup> See Lehrer, 1965, 1970, Lehrer and Paxson, 1968, Hilpinen, 1971, Johnsen 1974, Swain 1974, Barker 1976, Klein, 1971, 1976, 1980, Pollock, 1986, Moser 1989, Audi, 1993, de Almeida & Fett 2016. See Shope (1983) for a useful overview.

<sup>7</sup> For more on defeaters, see Pollock (1986), Piazza & Moretti (2018).

testimony that the lighting is non-standard. That gives me a defeater. But I might then learn that the testimony came from a pathological liar. That defeats the defeater and restores my *prima facie* reason.

Not all defeater-defeaters do this. Sometimes they give you new reasons rather than restoring the old ones.<sup>8</sup> For example, suppose I believe the wall is red because of my perceptual experience, then I'm told the lighting is non-standard but then I'm later told that the wall actually is red and the owners painted the wall red after installing the lighting just in case the bulb goes out so it will still look the same. The testimony about the lighting defeats my perceptual reason. The defeater is defeated by the information about the owners' back-up plan. That doesn't restore my perceptual reason: my vision still isn't to be trusted here. Instead, I now have a new reason to think the wall is red. When a defeater doesn't have a defeater-defeater that restores the subject's original justifying reason, I will call it a "non-restoring defeater".<sup>9</sup>

When the defeaters are part of the subject's possessed evidence, they are "psychological defeaters".<sup>10</sup> Sometimes facts of which the subject is unaware can be defeaters as well. Call these "factual defeaters".<sup>11</sup> This kind of defeater prevents a subject from having knowledge. However, it doesn't defeat the subject's justification. This is why Gettier cases involve justified, true belief with no knowledge. The important thing here is that non-restoring factual defeat is what precludes the subject's justifying reason from being related to the truth in the right way for them to have knowledge. If you have a true belief and your justifying reason for it has a non-restoring factual defeater, then the fact that you believe truly is due to knowledge-undermining luck.

Other kinds of luck don't involve a mismatch between justifying reasons and truth, however. For this reason, they don't preclude knowledge. Consider Nozick's (1981) case of the masked bandit. After robbing the bank, the bandit's mask slips for just a second. Coincidentally, you happen to be looking when it slips. You then acquire perceptual evidence that Jesse James is the bandit. You are lucky to have that evidence. It is a coincidence that you happened to be looking at the brief interval during which the mask was down. This luck doesn't undermine your knowledge, however. This is because the connection between your perceptual evidence and the belief it is evidence for (i.e., that Jesse James is the bandit) is still of the right sort. One way to flesh that idea out is that nothing defeats that perceptual reason. The mere fact that you could have easily failed to acquire the evidence is not a defeater. So, this kind of luck has no bearing on knowledge.

We can take this reasoning a step further. Not only can you be lucky to have the evidence you do; you can also be lucky not to have misleading evidence. Consider a variation of the Tom Grabit case from Lehrer and Paxson (1968). Suppose I am at the library and I see what looks like Tom Grabit stealing a library book. I then form the belief that Tom stole it. If I had left home just 5 min earlier, I would have

<sup>8</sup> This was first pointed out by Klein (1980), see also Pollock (1986: Appendix), De Almeida & Fett (2016).

<sup>9</sup> I get the term from de Almeida & Fett (2016). Klein (1980) uses similar terminology.

<sup>10</sup> Sometimes called "justificational defeat" (Steup, 1996).

<sup>11</sup> Sometimes called "propositional defeaters" Bergmann (2006), Piazza & Moretti (2018) or "knowledge defeaters" Audi (1993).

encountered Mrs. Grabit. She would have told me that Tom's identical twin brother, Tim is in town. She is telling everyone this. This is a lie she would have told to help cover up her son's burgeoning kleptomania, although I wouldn't have been in a position to discern this. Her testimony would have defeated my justification for believing that Tom stole the book. If I were to have heard the testimony, I would have had misleading evidence that would have prevented me from knowing that Tom stole the book. Luckily, I didn't hear it, so I do in fact know. Granted, the fact that this story is going around is a factual defeater. However, it is defeated by the further fact that she's lying. That restores my justifying reason rather than giving me new reasons. If I would have heard from Mrs. Grabit, however, I would have a psychological defeater. Nothing in my possessed evidence would defeat the defeater. So, I would have no justification and consequently no knowledge.

The defeasible reasoning theory has performed well so far. Some might still worry that it doesn't deal well with fake barn-style cases that feel like Gettier cases but involve non-inferential knowledge. Goldman (1976), for example, has said as much. For a response to this criticism, see Klein (1980).<sup>12</sup> Trying to settle this question would take us too far afield. Luckily, I will show that the defeasibility theory and its most formidable competitors both suit my purposes here.

Let us now consider the modal approach to making sense of knowledge undermining epistemic luck. Here are two modal properties that could potentially help shed light on knowledge undermining epistemic luck,

**Safety:** If an agent's belief that  $\Phi$  is safe, then, in nearly all (if not all) nearby possible worlds in which she forms her belief about  $\Phi$  the same way as she forms her belief in the actual world, that agent only believes that  $\Phi$  if  $\Phi$  is true.<sup>13</sup>

**Sensitivity:** If an agent's belief that  $\Phi$  is sensitive, then in the nearest possible worlds in which  $\Phi$  is false the agent does not form a belief that  $\Phi$  in the same way.<sup>14</sup>

Safety and sensitivity only have the potential to give us an account of knowledge undermining luck for beliefs in metaphysically contingent propositions. If the proposition believed is metaphysically necessary, then there are no worlds where it is false. So, we will restrict our attention for now to metaphysically contingent truths, although later we will consider another formulation of the safety principle meant to deal with necessary truths to see how it bears on my argument.

Some theorists use safety rather than sensitivity to explain knowledge and knowledge undermining luck (e.g., Hawthorne, 2004, Pritchard, 2005). Others use both

<sup>12</sup> For newer criticisms, see Feldman (2003), Foley (2012) and Turri (2012). In my view, an ample response to all three can be found in de Almeida & Fett (2016). Lasonen-Aarnio (2010a, b, 2014) and Baker-Hyatt and Benton (2015) raise a distinct set of worries I won't be able to address in this paper.

<sup>13</sup> This is a slightly modified version of the definition of safety from Pritchard (2005: 163).

<sup>14</sup> This is a modified version of Pritchard's (2005: 48) sensitivity principle that relativizes sensitivity to belief-forming methods.

criteria (e.g., Nozick 1981). In the cases we will be considering, the two will give the same verdict which will make it unnecessary to take a stance on this issue.

The intuitive idea the safety principle is supposed to capture is that when a belief is not safe, it could have easily been false. If a belief could have easily been false, then there is a sense in which the subject was lucky to get it right in the actual world. The intuitive idea the sensitivity principle is supposed to capture is that when a belief fails to be sensitive, the subject would have believed the same thing even if it were false. If a subject would have believed the same thing even if it were false, then her belief doesn't track the truth of the proposition believed. This might give one the impression that the subject was lucky, epistemically speaking, to be in one of the possible worlds where her belief is true. If she weren't in one of these worlds, she would be none the wiser.

These theories also have the potential to explain why some kinds of epistemic luck don't undermine knowledge and others do. Consider the Jesse James case from earlier. This kind of epistemic luck isn't intuitively knowledge undermining. The safety theory vindicates that intuition. Even though the subject was lucky to have the evidence they had and consequently to form their belief in the way they did, they could not have easily formed a false belief in that way. They had a good look at his face after all. The sensitivity theorist also has something to say here. In the nearest worlds in which it is false that the robber is Jesse James, either nobody robbed the bank or somebody else is riding away from the robbery. Either way, they won't believe it is Jesse James because they won't see someone who looks like him riding away from the bank.

Both theories also vindicate the intuitive result regarding the second kind of epistemic luck that does not undermine knowledge. Consider the Tom Grabit case again. The subject's belief is both safe and sensitive. It is safe assuming that the subject couldn't have easily mistaken someone else for Grabit. Since the identical twin is non-existent, this condition seems to be met. It is also sensitive for much the same reason. If Grabit did not steal the book and the subject formed a belief about whether he stole it by watching him, they would not have falsely believed that he stole it. Perhaps in nearby worlds in which the subject encounters Grabit's mother they would have formed a false belief that his identical twin stole it. This doesn't make their belief unsafe or insensitive, however, since they would have formed their belief in a different way (i.e., on the basis of testimony).

I will now go on to show that the second kind of luck I've identified can be caused by epistemic vice and still not undermine knowledge. To that end, I will take a brief detour and make a parallel point about resultant moral luck. Once we have done that and the epistemic analogue is in place, we will be in a position to see that there is a phenomenon of epistemically vicious knowledge and the consequences this has for a range of positions.

## 2 Resultant Moral Luck

In this section, I will discuss moral luck. I will do so because it is easier to see the general point I want to make in its moral form. Once we see the point, we will return to the epistemic case in the next section. I am doing this for the purpose of exposition. It is not important to my argument that moral luck exists. I am discussing it here because the reader is likely to be familiar with a certain way of thinking about moral luck. Adding one further wrinkle to moral luck, so understood, results in a moral version of the epistemic phenomenon that is my main concern here. So, even if this is the wrong way to think about moral luck, it might still help us understand a structurally similar phenomenon in epistemology.

A few kinds of moral luck have been identified. Here I am only interested in resultant moral luck. The basic idea here is that the moral standing of the agent is partially a function of the results of her actions. Worldly outcomes are not entirely within the agent's control, so her moral standing is also to some degree outside of her control. This means the agent might be morally worse because of bad luck or morally better because of good luck.

Here is an example of resultant moral luck. In *Manchester by the Sea*, an inebriated Lee Chandler leaves his home to get more beer. His family is at home asleep. He leaves a fire in the fireplace. His drunkenness prevents him from remembering to shut the screen to keep the embers in. When he returns home, the house has burned down with his children in it. The police interview him about the event the next day. The detective says something to the effect that plenty of people make a mistake and forget to shut the screen. It just so happens that Lee's mistake cost people their lives. It is obvious that Lee got unlucky. This is the detective's point. It is intuitive, albeit less obvious, to say that he had bad moral luck. That is, it is intuitive to say that Lee is morally worse off than the other people who forgot to shut their screens. Even if the other people were equally drunk and equally negligent, they are morally better off. Both Lee and these other people were culpably negligent. But Lee is also culpable for bringing about the deaths of several people. His counterparts aren't culpable for bringing about any deaths.

This is a controversial way to characterize the case. Some would rather give an alternative non-vindicating account of why we have different reactions when we consider Lee's case than we do his lucky counterparts.<sup>15</sup> They do this so that they can maintain that Lee and his counterparts are morally alike. They want to do this because they recoil from the idea that we can be blamed for things outside of our control. Lee and his counterparts exercised control in the same way. They performed the same intentional actions for the same reasons. So, they should receive the same moral assessment.

I can't give this line of thought the attention it deserves in this paper. I will just point out that, even if you doubt the existence of resultant moral luck, you might still agree with what I say in the epistemic case. Those that want to redescribe putative cases of resultant moral luck needn't also deny that there are Gettier cases, for

<sup>15</sup> See Richards (1986), Thomson (1989), Zimmerman (2002, 2006, 2015), Graham (2014), Rivera-López (2016), Khoury (2018). For recent opposition, see Hanna (2014).

example. Similarly, they could allow that there are kinds of epistemic luck that don't undermine knowledge. Even so, I will make my case by developing the moral analogue. I will do this because the reader is likely familiar with this way of thinking about resultant moral luck, even if they disagree with it. This will make it easier to understand the epistemic analogue in the next section.

Before going on to the epistemic analogue, one further piece needs to be put in place. We need to consider the possibility that moral luck can be explained by vice. Suppose you drive drunk and get lucky. You are lucky to achieve a moral standing I will call "in the clear". Being in the clear does not entail that you acted permissibly. You were negligent and negligence is not permissible. It means just that your negligence didn't result in any serious consequences this time. Being in the clear is the moral status the drunk driver hopes for when they get behind the wheel knowing it's a bad idea.

The fact that you are in the clear can be a consequence of your vice. Consider a drunk driver who would have hit a pedestrian had they not swerved at the last second. Their swerving had nothing to do with the fact that the pedestrian was there. They never saw the pedestrian and their reaction time had been slowed by drink too much for them to purposely get out of the way even if they had seen them. Rather, the driver swerved only because at the exact moment they needed to swerve, they saw a passing bicyclist and thought it would be funny to throw a beer can at them. Throwing the can caused them to jerk the steering wheel and this is what made them swerve out of the way. The driver is in the clear, but only because their vice manifested at the right time and in the right way.

### 3 Epistemically Vicious Knowledge

We can now return to the epistemic case.<sup>16</sup> We have already seen that a subject might be lucky to not have misleading evidence. This kind of luck doesn't undermine knowledge. I urge that, just as in the moral case, this kind of luck can be the result of epistemic vice and still not undermine knowledge. Just as your being in the clear can be explained by vice, so too can your not having misleading evidence. The fact that the driver is viciously chucking a beer can doesn't make them any less in the clear. Similarly, the fact that a subject is epistemically vicious doesn't make the luck that results from their vice knowledge-undermining. To sum up, the moral status of being in the clear is comparable to the status of knowledge. The former can be brought about by luck just as the latter can be brought about by the subject luckily not having misleading evidence. In either case, the relevant luck being brought about by a manifestation of vice doesn't change the status of the subject as either a knower or being in the clear. I will go on to apply the defeasible reasoning and modal traditions to the kind of case just discussed.

<sup>16</sup> For a very different way of seeing the moral luck/epistemic luck connection, see Slote (1985). He argues that whether one's belief in a future-tensed proposition at  $t$  is epistemically rational depends in some cases on whether that proposition turns out to be true at  $t+n$ . This seems false to me, though it is compatible with what I say here.



Before doing so, I should note that there may be differences between the two cases. The moral case involves the subject being lucky to be less blameworthy than their unlucky counterpart. The different epistemic statuses between the lucky and unlucky subjects might have nothing to do with blame but rather some other epistemic status. This doesn't change the general point that a subject could be lucky to achieve a positive standing of some sort even if that luck is explained by vice.

That is the main argument. I will now apply it to a case to see how it works in more detail. It is important to bear in mind that I am not making an argument primarily driven by intuitions about cases. If I were doing that, I would need to consider more of them. The theoretical considerations already introduced are meant to bear the burden of the argument. I am introducing a case here to illustrate the consequences of independently motivated theoretical considerations. With that in mind, let us consider an adaptation of Hilary Kornblith's (1983) case of the smug physicist.

**Smug Physicist:** A young physicist has a theory about the value of the gravitational constant. The theory is, let us suppose, well justified by her possessed evidence and true. She is so arrogant that when others try to tell her about defeaters that aren't part of her possessed evidence, she doesn't listen. That is, she simply doesn't encode what they are saying. She is too busy contemplating her own greatness to register their legitimate concerns. Now let us suppose that the defeaters are misleading, although nobody in the vignette has any reason to think so. The smug physicist's theory is correct and, furthermore, her justifying reasons are ultimately undefeated. That is, the defeater her colleagues are trying to tell her about has a defeater-defeater that restores her justifying reasons.

Let us begin assessing the smug physicist by the lights of the defeasible reasoning tradition. We should note that the subject does not have a psychological defeater. That is, her justifying reasons are undefeated by her total evidence. There is evidence she in some sense should have that defeats it, but she doesn't have that evidence. If she were to have even encoded the testimony of her colleagues, she would have a defeater in her possessed evidence.<sup>17</sup> Even if she didn't listen closely enough to know what they were saying specifically, she would at least have higher-order evidence so long as she encoded enough to register the disagreement. However, I am stipulating that she didn't even register that. So, she has a *prima facie* reason and no psychological defeaters. It is only because of her complacency and intellectual laziness that she doesn't have the misleading defeater. There is a factual defeater out there. This is the sort of thing that can undermine knowledge, but in this case it doesn't because it has a defeater-defeater that restores her justifying reason. This case is just like the Grabit variation from earlier, except Tom Grabit wasn't vicious.

Where does all this leave us? She has *prima facie* justification undefeated by her total evidence. Her justifying reason is also ultimately undefeated by the facts, since

<sup>17</sup> For this reason, there is no virtuous path to knowledge available to the agent. Although it doesn't matter to my main argument, I don't think this is essential to the case. Even if a virtuous path were available, so long as the subject took a vicious path instead, the explanation of how she in fact came to know would make essential reference to her vice. That said, the key point I want to make is most clear in cases where there is no virtuous path to knowledge available. Thanks to an anonymous reviewer for asking about this.

the only factual defeater is misleading and it has a defeater-defeater that restores her justifying reasons. So, she has a justified true belief and she is not Gettiered, at least according to the defeasible reasoning tradition. Gettier victims have factual defeaters and those factual defeaters are defeated, but in such a way that the Gettier victim's original justifying reasons are not restored.

Let us now move on to the modal properties. The subject's belief is in a nomological (but not metaphysical) necessity. So, it isn't false in any nearby worlds. It is also sensitive for a similar reason. The subject doesn't form a false belief using the same belief forming method in the nearest worlds in which the proposition is false because there is no life in those worlds. If the gravitational constant were slightly different, the universe would be unable to sustain life (assuming we hold fixed the values of several other constants).

One might object to the formulation of the safety principle I have adopted from Pritchard (2005). Perhaps the belief is unsafe, in the sense that counts, because the method by which the subject formed her belief could have easily produced a false belief in a different but relevantly similar proposition. This suggests an alternative safety principle,

**Safety II:** S's belief that *p* on the basis of *B* is safe if and only if: in most nearby worlds in which *S* forms a belief relevantly similar to *p* on the basis of *B*, *S*'s belief is true.

A proposal along these lines has been made in Pritchard's more recent work.<sup>18</sup> However, he has said little about how to individuate the bases of belief (cf. Hirvelä, 2019). In the present context, this matters greatly since whether the subject's belief satisfies **Safety II** depends on how we individuate her basis. On a very coarse-grained individuation, the subject based her belief on some combination let's say of induction and abduction on her possessed evidence, as scientist's typically do. Her evidence is not misleading and presumably competent induction and abduction are reliable enough to satisfy any reasonable standard. On an ultra-fine-grained individuation of her belief-forming process, she based her belief on induction and abduction on her possessed evidence plus tuning out the misleading testimony of her colleagues. Tuning out misleading testimony is certainly truth-conducive, so her belief satisfies **Safety II** on the ultra-fine-grained individuation of *B*. Her belief only fails to satisfy **Safety II** on an intermediately fine-grained individuation of *B* according to which she based her belief on induction and abduction on her possessed evidence plus tuning out the testimony of her colleagues. Perhaps some principled reason can be given for the intermediate individuation. Since most ways of individuating *B* have it that she satisfies **Safety II**, I will assume that she does and, at the very least, shift the burden of proof to those who disagree.

The upshot is that the subject's justifying reason relates to the truth non-accidentally, at least according to the familiar criteria. Her belief formed on the basis of that justifying reason is safe, sensitive and has no undefeated defeaters. I submit that the

<sup>18</sup> See his (2012: 256-7). In his (2007) he explicitly restricts the safety principle considered earlier in this paper to fully contingent propositions, which excludes nomological necessities.

subject believes truly not despite but largely because of her epistemic vice. It is only because of her complacency that she steered clear of misleading evidence. If she were to have had that evidence, she would have believed falsely that her hypothesis is false, suspended judgment or believed her hypothesis irrationally, depending on how we fill out the case. In none of these scenarios does she have knowledge.

It is worth considering her justificatory status for a moment. There are many theories of epistemic justification on the market, and it will be impossible to consider them individually here. Instead, I will focus on what seems like the most promising general strategy for denying that the subject has justification. Plausibly, sufficiently egregious irrationality can undermine justification. I will assume this going forward. This is not to assume that rationality is either necessary or sufficient for justification, only that egregious irrationality is incompatible with it. It is not to assume internalism. Some externalists give a reliabilist account of epistemic reasons that can be extended to cover epistemic rationality.<sup>19</sup> Others impose a “no-defeaters” condition on justification to prevent reliably produced but blatantly irrational beliefs from counting as justified.<sup>20</sup> At any rate, many can get on board with the idea that egregious irrationality can undermine justification. One might think that the subject is being sufficiently irrational in the case we are considering undermining her justification and, consequently, her knowledge.

Of all the criticisms that immediately spring to mind when we consider the smug physicist, irrationality is not obviously among them. She is smug, hubristic, complacent, irresponsible, intellectually lazy, etc. but none of this entails irrationality. My point is that the intuitive verdict about the smug physicist is that she exhibits all the vices mentioned above and maybe some others. Claiming that any of these vices entail irrationality is a substantive claim. To be sure, intellectual laziness (for instance) is not prudentially rational when your aim is to secure true beliefs. It isn't clear why this matters for justification though. The intellectually lazy know all sorts of things. The fact that they should have investigated more rigorously needn't prevent them from knowing them, at least so long as their beliefs are safe, sensitive, they have *prima facie* reasons and there aren't any non-restoring defeaters.

To sum up, the smug physicist doesn't seem to be irrational and her belief is non-accidentally true.<sup>21</sup> There is some luck involved, but not the knowledge undermining kind. I submit that we have a case of epistemically vicious knowledge. The subject knows not despite vice, but largely because of it.

<sup>19</sup> Beddor (2021), See also Comesaña (2010) for a reliabilist account of evidence.

<sup>20</sup> See Graham and Lyons (2021), who also give an externalist account of defeaters. See also Goldman (1979).

<sup>21</sup> I think we can give a similar analysis of Harman's (1980) Unopened Letter case, although it is more complicated and I don't have the space to pursue it here.

## 4 Ramifications

I now turn to the big picture implications of the conclusion reached in the last section. The first thing to note is that this creates difficulties for virtue-theoretic accounts of knowledge. According to these, a true belief is knowledge just in case the fact that the subject believed truly is explained (sufficiently or in the right way) by the subject's cognitive virtues. Different theorists have different things to say about what virtues are and how believing truly must be explained by virtue. I will consider a range of different proposals and identify difficulties for each of them.

I won't offer my own analysis of epistemic vice here. What I say here doesn't depend on any particular account of epistemic vice, only a sound pre-theoretical understanding of it that any well-developed theory ought to vindicate. Quassim Cassam proposes that epistemic vices are "character traits that impede effective and responsible inquiry" (2016). This analysis will suit my purposes, though my argument does not depend on it. Whatever epistemic vice ends up being, the smug physicist manifests it and her believing truly depends on this manifestation. Virtue epistemologists claim that knowledge is true belief that is creditable to the agent and it is creditable just in case it is explained (in the right way) by her virtues. The existence of vice raises the question of why her true belief is creditable to her if it depends just as much on vice as it does on virtue.

Before we consider some of the authors working in this tradition and how this problem bears on their specific views, we should preemptively respond to an objection I sometimes hear. A defender of virtue epistemology could claim that the unequal status of virtue and vice is to be explained by the fact that virtue bears on knowledge (or justification, for that matter) whereas vice does not.

This response presupposes the very thing that needs to be explained. The virtue theorist says that knowledge, over and above true belief, is a matter of the true belief being creditable to the agent. It is creditable just in case it is held because of a virtue. I am challenging the virtue theorist to explain why holding a true belief because of (plug in their preferred account of the "...because of..." relation here) a virtue is creditable to the agent even when it is also held because of (perhaps a different "...because of..." relation goes here) vice.

I am not going to provide an analysis of the dependence or "...because of..." relation that holds between the subject's true belief and her vice. It is clear that her believing truly depends on her vicious inquiry because that is what steered her clear of misleading evidence such that if she had it, she would not have believed truly. Unless given a principled reason to think otherwise, we should think that this dependence relation is just as explanatorily and epistemically significant as the ones that figure in virtue theoretic analyses of knowledge. This point can be appreciated prior to ironing out necessary and sufficient conditions for this dependence relation. In fact, it is compatible with that dependence relation being metaphysically and conceptually primitive, should that turn out to be the case.

Proponents of this order of explanation sometimes understand cognitive virtues as character virtues (i.e., roughly the same kind of thing as moral virtues such as humility and temperance). Linda Zagzebski, for example, holds this view and analyzes knowledge in terms of virtue here,

**Zagzebski:** Knowledge is a state of true belief arising out of acts of intellectual virtue (1996: 271).<sup>22</sup>

The fact that the agent holds the true belief out of intellectual virtue makes it creditable to her (Zagzebski, 2003). If you think of the connection between virtue and knowledge in this way, then you need to explain why intellectual vices can be so prominent in the explanation of why the subject believes truly without undermining knowledge. Even if virtue is necessary for knowledge, it isn't clear why the kind of vice at play in **smug physicist** isn't sufficient for undermining it. The virtue theorist owes us an explanation of this.

In more recent work, Linda Zagzebski has turned from analyzing knowledge in terms of character virtues to analyzing it in terms of "epistemic conscientiousness",

**New Zagzebski:** Knowledge is belief in which the believer gets to the truth because she acts in an epistemically conscientious way. (Zagzebski, 2008: 127)

Epistemic conscientiousness entails caring about truth, though not necessarily for its own sake and not necessarily in episodes of occurrent reflection (Zagzebski, 2008: 126; Zagzebski, 2014). The intellectual virtues are character traits one would try to develop if one were epistemically conscientious (Zagzebski, 2014: 143). This analysis faces a dilemma. Is the **smug physicist** epistemically conscientious? Affirmative and negative answers are both unpalatable because, intuitively, the subject's creditworthiness is diminished by her vice, but her knowledge isn't diminished at all.<sup>23</sup>

Let us first consider an affirmative answer. We can flesh it out as follows. Epistemic conscientiousness entails that one will *try* to cultivate the intellectual virtues. It does not entail complete success. Perhaps she has succeeded in developing her quantitative reasoning acumen but tried and failed to be more open-minded (though she isn't aware of her failure). If so, then she is epistemically conscientious despite her smugness. Furthermore, her conscientious quantitative reasoning figures prominently in the explanation of why she got to the truth. So, she knows. However, Zagzebski still owes us an explanation as to why the positive epistemic status associated with conscientiousness isn't canceled out by the negative epistemic status associated with her vice. This is especially important considering that her getting to the truth in this case depends on both. Recall that on Zagzebski's account, knowledge is true belief for which the subject deserves credit. It isn't clear why the dependence on vice doesn't undo the *pro tanto* creditworthiness that results from conscientiousness. Even if the subject is generally conscientious, this belief seems to be no less due to vice than conscientiousness. So why does she deserve credit for it?

Let us consider the other horn: a negative answer. If we go this way, the subject does not deserve credit and does not have knowledge, on her analysis. We should worry about denying knowledge here though considering that, as we saw in the last

<sup>22</sup> Zagzebski also offers two other definitions that she takes to be equivalent. See Kvanvig (2003) and Riggs (2009: 341) for similar views.

<sup>23</sup> Baehr (2006) and Lackey (2009) employ a similar argumentative tactic, but they focus on passivity rather than vice. Baehr (2014a) considers cases involving knowledge despite rather than because of vice.

section, the subject's belief is safe, sensitive and her justifying reasons are undefeated. According to these diagnostics, the physicist's belief is non-accidentally true. There is surely something wrong with the subject's epistemic conduct, but not anything that makes it an accident that she truly believes this proposition about the gravitational constant. When epistemic luck undermines knowledge, it is because of the relation between the subject's justifying reasons (or belief based on that evidence) and the truth of the proposition believed. According to the defeasible reasoning tradition, this is because the reasons it was based on are defeated by facts of which the subject was unaware. According to the modal tradition, it is because the belief, formed as it was, doesn't bear the right counterfactual relation to the truth of the proposition believed. Either way, the relation between her belief and the truth of the proposition believed seems to be of the right sort for knowledge.

At certain points Zagzebski acknowledges that there will be mixtures of vice and conscientiousness that are difficult for her view to assess (Zagzebski, 2014: 145). In response, she appeals to the vagueness of knowledge (Ibid). If the problem had to do with vagueness, we could solve it by stipulating a cut-off point. The problem is that, regardless of the cut-off point we choose, we will find ourselves on one or the other of the above horns. The subject is either sufficiently conscientious for knowledge, or she isn't.

Others think of virtues as, roughly, cognitive faculties such as memory, perception, etc. If you think of virtues this way, then the cognitive virtues that appear in your analysis of knowledge appear to be very different than the ethical virtues. Nonetheless, proponents of this approach sometimes argue that knowledge is true belief for which the subject deserves credit (e.g., Greco, 2003, 2010, 2012; Sosa, 2007: 22, Turri, 2011<sup>24</sup>). They get credit because they exercised their virtues. So, despite faculty virtues being different than character virtues, this order of explanation is still committed to saying that one becomes creditworthy by exercising them. This is why knowledge is epistemically better than mere true belief (Zagzebski, 1996, 2003; Riggs, 2009; Greco, 2010: Chap. 6).

We need an account then of why you deserve credit even if your vices are just as prominent in the explanation of why you believed truly as your virtues. In general, bringing about a good outcome by manifesting your vices is not something for which you deserve credit. Why are cases of epistemically vicious knowledge any different? Why don't the virtues and vices cancel each other out? Similarly, some working in this tradition say that knowledge must be due to virtue more than luck (Pritchard, 2012; Carter, 2016). If you think that, it seems no less reasonable to also say it must be due to virtue more than vice. Yet the smug physicist knows just as much because of vice as virtue. If she didn't inquire viciously, she would have had misleading evidence. Why then are virtue and vice equally explanatorily important but unequal in determining whether the subject deserves credit? Let us consider some specific proposals and see if they can give a satisfactory answer.

Turri (2011) offers the following,

**Turri:** Knowledge is true belief manifesting competence (i.e., faculty virtue).

<sup>24</sup> Pritchard (2007) agrees that the knowing subject deserves credit, though gives a different account of why.

Believing truly might be a manifestation of some faculty virtue (perhaps inductive or abductive reasoning) but be related in a different way to her character vice of intellectual smugness. Perhaps believing truly cannot itself be a manifestation of character vice, even if character vice is an enabling condition for it in some cases. If so, this might seem to get him out of trouble, since epistemically vicious knowledge doesn't depend on vice in the same way it depends on virtue.<sup>25</sup>

However, this should make us wonder why knowledge matters. If a subject believing truly depends on virtue in one way and vice in another and they are both interesting relations of dependence, then it isn't clear why the subject deserves credit for her true belief. It isn't clear why the dependence on vice doesn't cancel out whatever *pro tanto* creditworthiness comes from the dependence on virtue. Both kinds of dependence seem to be equally relevant to how we evaluate the subject, so it isn't clear why only one of them is relevant to whether she receives credit.

Now consider Sosa's (2007: 22) account,

**Sosa:** Knowledge is belief that is accurate because adroit.

Adroit beliefs issue from competences. How does the adroitness need to figure into the explanation of the accuracy? The adroitness of the belief needs to be an important part of the explanation of its truth. However, it does not need to be the most important. It can just be one of many important factors. This runs into the same problem: it leaves open the possibility that vice is no less explanatorily important than virtue. Why does virtue trump vice in such a case when the two are equally explanatorily important?

Greco (2003, 2010: 71) proposes,

**Greco:** S knows that p if and only if S believes the truth (with respect to p) because S's belief that p is produced by intellectual ability (i.e., faculty virtue).

The "...because..." relation is understood to be explanatory. In general, some parts of an explanation are more salient than others. Unlike Sosa's account, Greco's requires that the competences not just be a part of the explanation but that they be the most salient part of the explanation. Competences have default explanatory salience, but that default salience can be overridden by abnormalities of the sort that crop up in Gettier cases. Perhaps the presence of vice in the cases we have been considering is a similar abnormality. At first glance, this seems like an empirical claim of dubious merit considering the well-documented frequency of the base rate fallacy, affirming the consequent, confirmation bias, etc., all of which plausibly evince epistemic vices of some sort.<sup>26</sup> However, I will bracket this worry going forward as I think there are bigger problems for Greco.

Greco (2010: 74) claims that explanatory salience is a function of our practical interests and this is why faculty virtues enjoy default salience. However, it is unclear why the same wouldn't be true of vices. Just as good cognitive habits bear on our

<sup>25</sup> Whether this is true doesn't matter, I am just imagining a best-case scenario for Turri.

<sup>26</sup> See for example Kahneman et al. (1982).

practical interests, so do bad ones. Virtue and vice seem to be on equal footing in this respect. Perhaps this is why, intuitively, virtue and vice seem to be equally important in explaining why the subject believes truly in the cases we have been considering. If this is so, then the explanatory salience of virtue isn't demoted by an abnormality of greater explanatory salience, but rather supplemented by a vice of equal explanatory salience. Greco then faces a dilemma. He must either claim that when virtue and vice enjoy equal explanatory salience the tie goes to virtue and the agent gets credit; in which case he needs a principled reason for this claim. Or he can hold that vice must exceed the explanatory salience of virtue to undermine knowledge, in which case we have no explanation of why the agent deserves credit when virtue and vice are equally salient. We need a principled reason to think the tie goes to virtue so far as credit is concerned.

In more recent work, Greco (2012: 14) still claims that knowledge is true belief that is creditable to the agent, but now he says the following about when a true belief is creditable to the agent,

**New Greco:** A true belief is creditable to the agent just in case the subject's competences contribute to her believing truly in a way that serves the informational needs of the community.

This is not going to help in the case of the smug physicist precisely because *individual* scientists clinging to their theories can benefit the community even when those theories are not as well supported as rival theories (cf. Kitcher, 1993: 345ff). Even if all the available evidence favors vaccine A over vaccine B, it is still potentially good for the community to have some stubborn scientists championing vaccine B. This way the community as a whole has more comprehensive information about vaccines generally than they would if everyone just researched vaccine A. This is because, as Alfano & Levy put it, "Group deliberation harnesses the power of asymmetrical individual scrutiny" (2019: 907). That is, the labor of scrutiny is communally divided among biased individuals who will subject rival hypotheses to intense scrutiny. The group can benefit from having stubborn scientists among them, so long as the scientists stubbornly cling to different theories.

## 5 Conclusion

I hope to have convinced you that epistemically vicious knowledge is possible. That is, it is possible for an individual's knowledge to depend on her vices. This is possible for much the same reason that it is possible for a drunk driver's resultant moral luck to depend on some vice other than whatever vices got them behind the wheel in the first place. This presents complications for virtue-theoretic accounts of knowledge.

**Acknowledgements** I would like to thank Sanford Goldberg, Jennifer Lackey, Nate Lauffer, and Baron Reed for helpful comments.



## Declarations

**Conflict of Interests** I have no conflicting interests to declare.

## References

- Alfano, M., & Levy, N. (2019). Knowledge from Vice: A deeply social epistemology. *Mind*, 129(515), 887–915.
- Audi, R. (1993). *The structure of justification*. Cambridge University Press.
- Baehr, J. (2006). Character in Epistemology. *Philosophical Studies*, 128(3), 479–514.
- Baehr, J. (2014a). Are intellectually virtuous motives essential to knowledge? In M. Steup, J. Turri, & E. Sosa *Contemporary debates in epistemology* (2nd ed., pp. 133–40). Wiley-Blackwell
- Baehr, J. (2014b). Reply to Zagzebski. In M. Steup, J. Turri, E. Sosa, (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 146–149). Wiley-Blackwell
- Baker-Hytch, M., & Benton, M. (2015). Defeatism defeated. *Philosophical Perspectives*, 29(1), 40–66.
- Bedder, B. (2021). Reasons for Reliabilism. In J. Brown, & M. Simion (Eds.), *Reasons, justification and defeat* (pp. 146–177). Oxford University Press.
- Barker, J. (1976). What you don't know won't hurt you? *American Philosophical Quarterly*, 13(4), 303–308.
- Bergmann, M. (2006). *Justification without awareness*. Oxford University Press.
- Carter, J. A. (2016). Robust virtue epistemology as anti-luck epistemology: a new solution. *Pacific Philosophical Quarterly*, 97(1), 140–155.
- Cassam, Q. (2016). Vice Epistemology. *The Monist*, 99(2), 159–180.
- Comesaña, J. (2010). Evidentialist reliabilism. *Nous*, 44(4), 571–600.
- de Almeida, C. & Fett, J.R. (2016). Defeasibility & gettierization: a reminder. *Australasian Journal of Philosophy*, 94(1), 152–169.
- Feldman, R. (2003). Epistemology. *Tijdschrift Voor Filosofie*, 68(2), 429–429.
- Foley, R. (2012). *When is true belief knowledge?* Princeton University Press.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Goldman, A. (1979). What is justified belief? In George S. Pappas (Ed.), *Justification and knowledge: new studies in epistemology* (pp. 1–25). Reidel. reprinted in his (2012). *Reliabilism and contemporary epistemology* (pp. 29–49). Oxford University Press.
- Graham, P. (2014). A sketch of a theory of Blameworthiness. *Philosophy & Phenomenological Research*, 88, 388–409.
- Graham, P., & Lyons, J. (2021). The structure of defeat: Pollock's Evidentialism, Lackey's Framework, and prospects for Reliabilism. In J. Brown, & M. Simion (Eds.), *Reasons, justification and defeat* (pp. 39–68). Oxford University Press.
- Greco, J. (2003). Knowledge as credit for true belief. In M. DePaul, & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Clarendon: 111–34.
- Greco, J. (2010). *Achieving knowledge: A Virtue Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Greco, J. (2012). A (different) Virtue Epistemology. *Philosophy and Phenomenological Research Vol*, 85(1), 1–26.
- Hanna, N. (2014). *Moral Luck defended*. *Nous*: 683–98.
- Harman, G. (1980). Reasoning and explanatory coherence. *American Philosophical Quarterly*, 17(2), 151–157.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.
- Hilpinen, R. (1971). Knowledge & Justification. *Ajatus*, 33, 7–39.
- Hirvelä, J. (2019). Global safety: How to deal with necessary truths. *Synthese*, 196(3), 1167–1186.
- Johnsen, B. (1974). Knowledge. *Philosophical Studies*, 25(4), 274–282.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.
- Khoury, A. (2018). The objects of Moral responsibility. *Philosophical Studies*, 175, 1357–1381.
- Kitcher, P. (1993). *The Advancement of Science*. Oxford University Press.

- Klein, P. (1971). A proposed definition of Propositional Knowledge. *Journal of Philosophy*, 68, 471–482.
- Klein, P. (1976). Knowledge, causality, Indefeasibility. *Journal of Philosophy*, 76, 792–812.
- Klein, P. (1980). Misleading evidence and the restoration of justification. *Philosophical Studies*, 37, 81–89.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *Philosophical Review*, XCII(1983), 33–48.
- Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies*, 142(1), 27–42.
- Lasonen-Aarnio, M. (2010a). Is there a viable account of well-founded belief. *Erkenntnis*, 72(2), 205–231.
- Lasonen-Aarnio, M. (2010b). Unreasonable knowledge. *Philosophical Perspectives*, 24(1), 1–21.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88.2, 314–345.
- Lehrer, K. (1965). Knowledge, Truth and evidence. *Analysis*, 25, 168–175.
- Lehrer, K. (1970). The Fourth Condition for knowledge: A defense. *Review of Metaphysics*, 24, 122–128.
- Lehrer, K., & Paxson, T. (1968). Knowledge: Undefeated justified true belief. *Journal of Philosophy*, 66, 225–237.
- Moretti, L., & Piazza, T. (2018). Defeaters in Epistemology: Introduction to the Special Issue. *Synthese*, 195, 2845–2854.
- Moser, P. (1989). *Knowledge and Evidence*. Cambridge University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.
- Olsson, E. (2015). Gettier and the Method of Explication: A 60 year-old solution to a 50 Year Old Problem. *Philosophical Studies*, 172, 57–72.
- Pollock, J. (1986). *Contemporary Theories of Knowledge* 1st Edition. Rowman & Littlefield: Totowa, NJ.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford University Press.
- Pritchard, D. (2007). Anti-luck Epistemology. *Synthese*, 158(3), 277–297.
- Pritchard, D. (2012). Anti-luck Virtue Epistemology. *The Journal of Philosophy Vol.*, 109 No. 3: 247–279.
- Richards, N. (1986). Luck and Desert. *Mind*, 95, 198–209.
- Riggs, W. (2009). *Epistemic Value*. Oxford University Press.
- Rivera-López, E. (2016). How to reject resultant Moral luck alone. *Journal of Value Inquiry*, 50, 415–423.
- Shope, R. (1983). *The analysis of knowing: A decade of Research*. Princeton University Press.
- Slote, M. (1985). *Common-sense morality and consequentialism*. Routledge & Kegan Paul. Boston.
- Sosa, E. (2007). *A Virtue Epistemology: Apt belief and reflective knowledge, volume 1*. Clarendon.
- Steup, M. (1996). *An introduction to contemporary epistemology*. Prentice Hall.
- Swain, M. (1974). Epistemic defeasibility. *American Philosophical Quarterly*, 11(1), 15–25.
- Thomson, J. (1989). Morality and bad luck. *Metaphilosophy*, 20(3/4), 203–221.
- Turri, J. (2011). Manifest failure: The Gettier Problem Solved. *Philosophers' Imprint*, 11, 1–11.
- Turri, J. (2012). In Gettier's wake. In S. Hetherington (Ed.), *Epistemology: The key thinkers* (pp. 214–229). Continuum.
- Weatherston, B. (2003). What good are counterexamples? *Philosophical Studies*, 115, 1–31.
- Weinberg, J. (2017). Knowledge, noise and curve-fitting: a methodological argument for justified true belief? In R. Borges, C. de Almeida, & P. Klein, (Eds.), *Explaining knowledge: new essays on the gettier problem*, (pp. 253–272). Oxford University Press
- Zagzebski, L. (1996). *Virtues of the mind: An Inquiry into the nature of Virtue and the ethical foundations of knowledge*. Cambridge University Press.
- Zagzebski, L. (2003). The search for the source of the Epistemic Good. *Metaphilosophy*, 34(1–2), 12–28.
- Zagzebski, L. (2008). *On Epistemology*. Wadsworth: Belmont, CA.
- Zagzebski, L. (2014). Knowledge and the motive for truth. In M. Steup, J. Turri, & E. Sosa, (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 140–146). Wiley-Blackwell
- Zimmerman, M. (2002). Taking luck seriously. *The Journal of Philosophy*, 99, 553–576.
- Zimmerman, M. (2006). Moral Luck: A partial map. *Canadian Journal of Philosophy*, 36, 585–608.
- Zimmerman, M. (2015). Moral Luck re-examined. *Oxford Studies in Agency & Responsibility*, 3, 136–158.

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.