



Good reasons are apparent to the knowing subject

Spencer Paulson^{1,2} 

Received: 29 November 2022 / Accepted: 22 June 2023
© The Author(s), under exclusive licence to Springer Nature B.V. 2023

Abstract

Reasons rationalize beliefs. Reasons, when all goes well, turn true beliefs into knowledge. I am interested in the relationship between these aspects of reasons. Without a proper understanding of their relationship, the theory of knowledge will be less illuminating than it ought to be. I hope to show that previous accounts have failed to account for this relationship. This has resulted in a tendency to focus on justification rather than knowledge. It has also resulted in many becoming skeptical about the prospects for an analysis of knowledge. The skepticism is misplaced and the tendency can be fixed without sacrificing any insights. The solution is to see how good reasons (in a sense to be articulated) are apparent (in a sense to be articulated) to the knowing subject. Once this claim is unpacked, we see that it is an illuminating analysis of knowledge in terms of distinct but intimately related aspects of epistemic assessment. It helps us see the value of knowledge (over and above true belief) and it affords us a unified, “reasons-first” metaepistemology. The resulting picture is neither a familiar kind of internalism or externalism.

Keywords Reasons-first · Analysis of knowledge · Gettier problem · Metaepistemology

I’m interested in making sense of the following two claims:

- (1) The subject’s reasons rationalize her belief.
- (2) Those same reasons, when all goes well, turn her true belief into knowledge.

These claims both seem very plausible, yet it is difficult to adequately account for them jointly. The problem is that (1) seems to require that we understand reasons in such a way that it is difficult to make sense of (2). If we understand reasons¹ in such

¹ Before getting started I should mention that I like to put my point in terms of reasons, though for my purposes it would also be ok to put the points I make in terms of evidence.

✉ Spencer Paulson
spencerpaulson2023@u.northwestern.edu

¹ Northwestern University, Evanston, USA

² Evanston, USA

a way that their essential property is to rationalize belief, then it seemingly must be an accidental property of theirs that turns true belief into knowledge. This is because Gettier (1963) has shown us that rationality doesn't turn true belief into knowledge. Our interest in knowledge is an interest in the subject's reasons, though it isn't just an interest in whether they are good reasons. It is an interest in whether they are good reasons that stand in the right relation to the facts: something ostensibly foreign to their role of rationalizing belief. To be clear, I am not saying that theorists intend for these roles to be foreign to each other. Rather, I am saying (and will argue below) that we haven't yet seen an account that vindicates our sense that they are not foreign to each other.

This is unfortunate. Knowledge seems too much like the intersection of unrelated concerns to be a philosophically interesting phenomenon. At the very least, it would be better to have an account of how these ostensibly unrelated aspects of the subject's reasons are really distinct aspects of a single concern. Otherwise, the resulting account is in danger of sounding ad hoc (cf., Kvanvig, 2003; Pollock, 1986, p. 183; Schroeder, 2015; Williamson, 2000, p. 31. It would lack the elegance and theoretical illumination we should hope for (cf. Zagzebski, 2017). If knowledge is merely the intersection of otherwise unrelated properties, it is unclear why we care so much about that intersection (cf. Williamson, 2000, p. 31).

I will argue that distinguishing subjective and objective senses of "good reason" to tackle (1) and (2) (respectively) and requiring that the knower's reasons be both is not enough. Very briefly, the subjective sense of "good reason" or "normative reason" has to do with rationality and the objective sense has to do with some kind of cosmic serendipity. These seem unrelated to one another and simply requiring that they coincide doesn't address that problem. Why is it those two things, rather than two other things, that must coincide? I will argue that a similar worry arises for attempts in the defeasible reasoning literature according to which (1) is a matter of justification and (2) is a matter of that justification being "robust" (Schroeder, 2015, "defeat-proof" (de Almeida & Fett, 2016; Klein, 2017) or "truth-resistant" (Moser, 1988, 1989).

The key is to understand the "cosmic serendipity" aspect of knowledge in such a way that it is distinct from but not unrelated to the role that reasons play in rationalizing the subject's belief. This will give us an illuminating account of how the two aspects of the subject's reasons that interest us hang together. I will end by giving an account that does this by showing how the sense of "good reason" required to make sense of (1) is related to the sense required to make sense of (2). More specifically, a reason that is good in the former sense is apparently good in the second. Once this idea is unpacked, we will see that our interest in knowledge is an interest in whether the subject's reasons really are as they appear. Good reasons (in a sense to be articulated) are apparent (in a sense to be articulated) to the knowing subject.

In Sect. 1, I will briefly discuss claims (1) and (2). In Sect. 2, I will discuss how our efforts to make sense of (1) seem to put us in a difficult position with respect to (2). In Sect. 3, I will critically assess what I take to be the best attempts to address this difficulty. In Sect. 4, I will present my own account of how our interests in (1) and (2) relate and the picture of knowledge that emerges.

1 Reasons in epistemology

The first of the two claims that I want to consider is that the subject's reasons rationalize her belief. By "rationalize" people sometimes mean a post hoc confabulation. This is not what I have in mind. Rather, I mean that the subject's reasons make it the case that her belief is rational, at least in the case in which her belief is rational. Basic or foundational beliefs present a complication here. I will discuss that in Sect. 4.² In any event, what draws many of us into epistemology is that we wonder whether we are as rational as we take ourselves to be (cf. Audi, 1993; BonJour, 2003, p. 40, 2010; Fumerton, 1995; Smithies, 2015). We want to know if our beliefs are rational and if the question of rationality doesn't arise for some beliefs, then we want to know what is special about them.

When we are looking into the rationality of a subject's belief, we look at how reasonable that belief is. If her reasons make her belief reasonable it is because of the state of her total evidence. I am having a perceptual experience as of a Kentia Palm. Is that a good reason for me to believe that there is one where I am looking? Relative to the total evidence I in fact have, it is. However, if it were part of my evidence that I am somewhere densely populated with artificial Kentia replicas, then it wouldn't be such a good reason. Put another way, relative to the evidence I in fact possess, I have sufficient subjective normative reason to believe that there is a Kentia Palm where I am looking. If it were part of my evidence that there are many artificial replicas nearby, then I would not have sufficient subjective normative reason to believe it. "Subjective" here simply means that the normative reason is sufficient insofar as it is being assessed relative to the subject's total evidence.³ I use "good reason" and "sufficient reason" interchangeably in this paper. I do this for stylistic reasons. At certain points "sufficient" would be clunky while "good" sounds fine.

The above is a very natural line of thought and there is surely something right about it. The basic idea is that we start with the idea that reasons rationalize beliefs. Our recognition that they play this role informs our understanding of what a good reason is. I take this idea about the connection between reasons and rationality to be a well-ensconced part of common-sense that is worth trying to preserve.

Epistemology is also about knowledge. How does this relate to our interest in rationalization? At first glance, there seems to be a very close relation. There are a few ways to see this point. When we ask someone how they know something, we are asking for their reasons (Pollock, 1986, p. 5; Scanlon, 1998). This suggests (2). If we answer that question by giving our reasons, this suggests that when our true beliefs are knowledge, it is because of our reasons. A related point comes up when we try to explain why more than mere true belief is needed for knowledge. Someone could acquire a true belief by making a lucky guess. Surely this isn't knowledge. What is missing? The obvious answer is that they didn't have any good reason for believing as they did (cf. BonJour, 2010, Chapter 2; Steup, 1996, Chapter 1).

² Non-foundationalists say that reasonableness always requires having further reasons. Foundationalists deny this. I will remain neutral until later in the paper.

³ As opposed to the facts, in which case we would be considering its sufficiency as an objective normative reason. More on objective normative reasons and their relation to subjective normative reasons below.

If this much were right, then we would have an adequate account of (1) and (2). In order to make sense of (1) we come to believe that a good reason is one that rationalizes the subject's belief. A good reason, so understood, is also what turns a true belief into knowledge. Our interest in knowledge over and above true belief is an interest in whether the subject has good reasons. It seems hardly surprising that we would be interested in knowledge, if that is what the interest comes to. Offering and assessing reasons is an important aspect, perhaps the most important aspect, of our intellectual lives.⁴ Our capacity for this kind of assessment makes us persons.

What I am doing here is taking our practices at face value and then trying to see if we can vindicate them (cf. BonJour, 2003, p. 40 for a similar view about the task of epistemology). When we ask how someone knows, we ask for her reasons. The assessment of reasons is second nature to us. We understand this kind of assessment in all its nuance, at least implicitly, very well. If we can understand both (1) and (2) in terms of a single kind of assessment that is second nature to us, then we have shed light on (1) and (2) in such a way as to vindicate our practice at least in the following way. Our way of inquiring into whether someone knows presupposes that her knowledge is a function of her reasons. Our appreciation of reasons, our rationality, is a deep feature of how we understand ourselves. If we understand (2) in terms of the role reasons play in making (1) true, then we understand our interest in knowledge in terms of an aspect of ourselves that we take to be very important to who we are. The centrality of knowledge to our social commerce and private inquiries is something that makes sense to us and that we endorse once we see what knowledge is. Once we understand it this way, we see that it makes sense that we care so much about it. This all pulls in the direction of a "reasons-first" epistemology (see Kvanvig, 2017; Schroeder, 2015, 2021). We get what we want by making sense of knowledge in terms of our prior understanding of reasons and their assessment.⁵

2 Hybrid knowledge

The problem with the picture that emerged at the end of the last section is that good reasons, so understood, are not what turn true belief into knowledge. This is the upshot of the Gettier (1963) problem. The subject doesn't just need good reasons, the world needs to make her belief true in a way that lines up with those reasons. This is all purposely vague, since it is an attempt to state the problem without presupposing any particular solution to it. After setting up the problem in its full generality, I will argue in the next section that the most promising solution on the market comes up short.

Our interest in knowledge over and above true belief still seems to be an interest in the subject's reasons. The Gettier problem is a problem with the subject's reasons not lining up with the facts in the right way. I will continue to assume that rationality is required for knowledge (which will make my conclusions in this paper conditional on

⁴ For various takes on this idea see Sellars (1956/1997), BonJour (1985, p. 7), Korsgaard (1996, Lecture 3), Scanlon (1998, Introduction), Moran (2001, Chapter 2), Boyle (2012), Smithies (2016). For a dissenting view, see Kornblith (2011).

⁵ Some readers may worry that this is inconsistent with our tendency to attribute knowledge to small children and animals. I disagree, although the point will have to be made in a different paper.

that assumption). So, for the subject to know, she must have good reasons (i.e., ones that rationalize her belief). In addition to being good, those reasons need to harmonize with the facts in the right way.

So, our interest in knowledge is not just an interest in whether the subject's reasons are good. It is an interest in whether they are good reasons with some further property that is orthogonal to their goodness as reasons. In other words, it is not merely an interest in their essential properties but their accidental ones as well. This point does not depend on the ontology of reasons. When I talk about their essential properties here, I simply mean how they perform with respect to the role they are supposed to play as reasons.

A difficulty emerges here. Why is there a single phenomenon with both essential and accidental properties of reasons as components? Put another way, why do the goodness of the subject's reasons and this other property of theirs compose a single, unified phenomenon: knowledge? If this other property of reasons is another sense of "good reason", the problem rears its head in a slightly different form: we have two lexical entries for "good reason", how do they relate? Without an answer, I worry that the concept of knowledge appears to be gerrymandered. An object is gerrymandered just in case it is the mere sum of distinct objects. The objects countenanced by mereological universalists are like this. My hiking shoes and the Queen's eldest Corgi have no interesting relationship. There is nothing in virtue of which they belong together. So, the "object" that is their sum is gerrymandered. Something similar can apply to composite concepts. If their components aren't interestingly related, they are gerrymandered.

My worry is that a certain way of distinguishing the components of knowledge makes it appear to be one of these concepts. That way of distinguishing them seems to be inevitable. Rational subjects have good reasons. Otherwise, they wouldn't be rational. Yet this leaves open the possibility that their good reasons won't line up with the facts in the right way for knowledge. The problem is that it is hard to see how this further requirement of "lining up" relates to rationality.

The interest in rationality seems to be self-contained. It doesn't seem to be a component part of a larger interest. It is a matter of how the evidence within the subject's perspective functions and doesn't seem to bear any interesting relationship to how those reasons function relative to facts beyond the subject's ken. The tendency to give a theory of justification/rationality that abstracts away from how the subject's reasons line up with the facts makes it seem, perhaps contrary to the intentions of the theorists doing this, that we can understand each without reference to the other (more on this below). So, although the anti-Gettier condition is supposed to be the connection between justification/rationality and truth,⁶ the connection is starting to

⁶ Perhaps better, a (rather than the) connection between justification and truth. Justification is already connected to truth in the minimal sense that your justification is something that makes it rational for you to regard something as true. The problem is that your justification needs another connection to truth for you to have knowledge; this is the upshot of the Gettier problem. Another way of voicing the concern is that we have two different truth-connections and they seem to both be interesting but not like two different aspects of a single interesting thing.

look like a relation between largely independent concerns (i.e., what the subject's evidence makes it reasonable for her to believe, on the one hand, and how many important worldly facts she happens to be aware of on the other).

We should try to state what it is that concerns us when we inquire into knowledge in such a way that we can then see both aspects of the subject's reasons as distinct but complementary aspects of it. Saying they are interestingly related because they are components of knowledge presupposes what we would like to have vindicated: that knowledge is a unified, philosophically interesting phenomenon. To be clear, it isn't surprising that we sometimes find gerrymandered properties practically interesting. For practical purposes it sometimes makes sense to fixate on a cluster of essential and accidental properties (e.g., the previous owner of a baseball bat). But it would be strange to find such a motley collection philosophically interesting enough center an entire discipline on the study of it.

It is possible that we will never be able to vindicate our sense that the components of knowledge are interestingly related. Perhaps our inquiry into knowledge will embarrass us just as our inquiry into Jade did.⁷ If this is where we end up, then we know which properties intersect to form knowledge, but it becomes mysterious why we take such philosophical interest in that intersection (cf. Williamson, 2000, p. 31 for a related point⁸). If our account is a mere aggregate of ingredients, then it lacks depth (even if the aggregated properties are individually deep).⁹ The list of ingredients seems ad hoc until we give a principled account of why they belong together (Ibid, cf. Kvanvig, 2003; Pollock, 1986, p. 183; Schroeder, 2015). Our account of the phenomenon lacks the elegance and theoretical illumination we should want (cf. Zagzebski, 2017). Some phenomena lack depth. It is conceivable that knowledge is one of these, though that would be surprising and a bit of a letdown.

Compare: good deeds are an interesting subject matter. Recognition is an interesting subject matter. It doesn't follow that good deeds that get recognized form a similarly interesting or philosophically important subject matter. If we had a branch of philosophy that concerned itself primarily with good deeds that get recognized, that would be surprising. We would expect that philosophers would either be able to state an interesting theoretical connection between the goodness of a deed and its recognition, or the discipline would shift its focus. In the latter case, philosophers would say that the philosophically interesting thing is either the goodness of the deed or whether it gets recognized. We should expect the former. It seems like an essential property of the deed: it has to do with the excellence of the deed considered as a deed. The latter is an accidental relational property of the deed. It isn't clear why I should be interested in it at all insofar as I am interested in deeds.

Something similar happened in epistemology. There is a temptation to think of (1) as giving us the essential properties of reasons, in the sense that a reason is a good one just in case it rationalizes the subject's belief. There is a clear sense, after all, in

⁷ Cf. Putnam (1975).

⁸ Kirkham (1984) embraces the consequence.

⁹ One might object that the anti-Gettier condition is meant to be the connection between justification and truth, so it isn't a mere aggregate. My concern is that the connection just seems like another dimension of assessment entirely. I take it this is the worry that the comments about JTB + analyses being ad hoc and unilluminating are gesturing toward.

which Gettier victims have good reasons for what they believe, despite their failure to have knowledge. (2) then appears to be an accidental relational property of the reason. This makes a phenomenon composed of (1) and (2) seem like an odd thing to have as the primary subject matter of an entire discipline. So, there is a tendency to say that epistemology is really about (1) and (2) is mentioned, if at all, as a codicil.

In *The Structure of Empirical Justification*, BonJour is giving a theory of the role reasons play in rationalizing belief. He briefly notes that we need more than truth and rationality for knowledge, but he brackets the point because the relation of this something more to justification/rationality is “obscure” to him (BonJour, 1985, p. 5). So, he takes the theory of justification to be “the most fundamental sort of epistemological theorizing” (BonJour, 2003, p. 40), the one that is interesting from a first-person perspective (BonJour, 2003, p. 39). Fumerton also takes rationality to be the philosophically interesting (from a first-person perspective) “core of a concept of knowledge” (1995, 5ff, cf. Kaplan, 1985 for the same argument), which makes the philosophically less interesting periphery seem like an afterthought. Pollock and Cruz go further. Not only do they take (1) to be the philosophically interesting component of knowledge, they think that (2) makes knowledge a perverse concept:

What the Gettier problem really shows is what a perverse concept knowledge is. One can do everything with complete epistemic propriety, and be right, and yet lack knowledge because of some accident about the way the world is. Why do we employ such a concept? (Pollock & Cruz, 1999, p. 14)

This is given as a reason for focusing their epistemology book (curiously titled *Contemporary Theories of Knowledge*) on justification rather than knowledge. Epistemology is to be done from the first-person perspective and only the theory of justification is action-guiding (Pollock & Cruz, 1999, p. 124). I submit that when so many of the leading figures in the theory of knowledge tell us that only one component of it is *the* philosophically interesting part, something has gone wrong. Yet this is what we should expect if the account they offer is gerrymandered. This leads me to think that we haven’t really seen our way to clarity as to how the anti-Gettier condition is the *connection* between rationality/justification and truth. Despite knowing on some level that we can only understand what rationality is in terms of this connection, there is a tendency to abstract away from that relation when theorizing about it. We then end up (implicitly at least) thinking of it as a self-contained phenomenon and then face the difficult task of trying to glue it together with another self-contained phenomenon later.

3 Two kinds of good reason

The basic idea I want to consider here is that we can make sense of the role reasons play in converting true belief into knowledge in terms of the role they play in rationalizing belief. We can do this by requiring that when they convert true belief into knowledge, they rationalize the subject’s belief in a way that does not depend on her ignorance of the facts. This seems like a step in the right direction.

This is the basic idea behind defeasibility solutions to the Gettier problem.¹⁰ Robert Audi puts the point well here,

[The conditions I propose are] are related to, and to some extent unified by, justification in the following way: if [the subject] *believed* that his circumstances were as they are in relation to the beliefs in question, he would not be justified in holding them. More generally, it is at least often the case that what undermines S's would-be knowledge without undermining his justification *would* undermine his justification if he believed it. And this makes it natural to say, though it does not require us to say, that much defeated knowledge is defeated because S's justification is in some way insecure or defective. (Audi, 1993, p. 205)

In the chapter from which the quote is taken, Audi is talking about Gettier-style cases having to do with perceptual knowledge specifically. However, the same point can be made in the classic Gettier cases as well.

Audi puts the point in terms of justification, but it can easily be translated back into the terms I have used to frame the paper. Internalist justification is a matter of having sufficiently good reasons or adequate evidence. This comes to the same thing as the subject's belief being rationalized. Of course, a belief can be both rationalized and true yet fail to be knowledge. What goes wrong in such a case? Audi's point is that, at least often, the problem is that the reasons that rationalize the subject's belief are defective or insecure. Their insecurity consists in the fact that, were the subject to believe that her circumstances are as they really are, that reason would no longer be able to play its rationalizing role for her.

Essentially the same point is made by a number of other authors. Sometimes instead of talking about a subject's justifying reasons being defective or insecure they say that in the kind of case Audi is imagining those reasons are "rotten" (de Almeida & Fett, 2016), "not worth having" (Klein, 2017), or they don't "epistemize" (Steup, 1996, Chapter 1). Similarly, when the subject knows her justifying reasons are "robust" (Schroeder, 2015), "truth-resistant" (Moser, 1988, 1989) or "defeat-proof" (Klein, 2017). The important point for our purposes is how the reason's merits or defects are understood in terms of whether it could rationalize the belief if she were to know the facts about her situation.

Suppose a student tells the professor that he owns a Ferrari, and the professor forms the belief that someone in her class owns a Ferrari (cf. Lehrer, 1965). It turns out that someone in her class does own a Ferrari, but it is not the student who told her he did. That student is a pathological liar, although the professor had no reason to suspect this. The professor has a true belief that is rationalized by the testimony she has received. It nonetheless falls short of knowledge because it wouldn't rationalize her belief if she were aware that he is a pathological liar.

Now we have an explanatory connection between the role that reasons play in converting true belief into knowledge and the role they play in rationalizing beliefs. This makes some progress toward addressing the gerrymandering worry, since the absence of any kind of interesting explanatory connections between the components

¹⁰ 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 and Fett (2016). See Shope (1983) for a useful overview.

of a theory are what makes the theory gerrymandered. Although this is a good start, I worry that more needs to be said.

First, the point is made as a counterfactual. If learning the facts about her situation would prevent the reason that rationalizes her belief from continuing to do so, then her reason does not turn her true belief into knowledge. Why does this counterfactual bear on her *actual* epistemic status? We can all agree that her reason *wouldn't* rationalize her belief if she were to know more. But she doesn't know more. Why is it the case that her reasons *are* defective? Why isn't it just that they would be if she were to know more? After all, as things stand, they rationalize her belief. Indeed, it is a little strange to say that her justification is defective, rotten, etc. Considered as justification, it seems fine. She is just as justified as the knowing subject, otherwise Gettier cases wouldn't be counterexamples to JTB analyses. The same point can be translated into reasons-talk. If we understand what it is to be a good reason the way we have been, then her reasons aren't defective, rotten, etc. They do exactly what they are supposed to do: rationalize her belief relative to the total evidence she in fact has. So, although there is an explanatory connection between (1) and (2), it isn't of the right sort.¹¹

There are other problems with putting the point in counterfactual terms. There are many ways to make the antecedent of a counterfactual true and some of them have unintended consequences (for an extended discussion of this point in relation to the Gettier problem, see Shope, 1983). For this reason, we should make essentially the same point slightly differently. Instead of stating the point in terms of whether the subject's reasons would rationalize her belief were she to know the facts, they state it in terms of whether there are facts that either trump or mitigate her reasons. The key idea here is that in order for the subject to have knowledge, her justifying reasons must be good in both the subjective normative and the objective normative sense. That is, her reasons must be good assessed relative to her total evidence and also good assessed relative to the facts. In Lehrer's Ferrari example, the subject's reasons are good insofar as we assess them relative to her total evidence (i.e., she has a subjective normative reason). However, they are not good insofar as we assess them relative to the facts (i.e., she does not have an objective normative reason). Her testimonial reason is mitigated (or undercut) by the fact that the testimony came from a pathological liar.

What I say here won't depend on any particular account of the necessary and sufficient conditions for reasons being good in either sense. There is a cluster of very similar proposals in the defeasible reasoning tradition.¹² More recently, Horty (2007, 2012, cf. Schroeder, 2021, 92 ff.) has developed a formal model of defeasible reasoning using default logic. Similarly, Schroeder (2007, cf. 2015) has given a recursive specification of how reasons are weighted that can be applied to either subjective normative or objective normative reasons. The differences between these views needn't concern us here. I am interested in the genus of which they are species.

Each species gives us an account of which conditions must obtain for a reason to be good in the subjective normative and the objective normative sense. Even after we

¹¹ It is sometimes argued that knowledge is more explanatorily powerful than mere rational true belief because it is more robust in the face of new evidence and hence less likely to be lost after one learns more (Williamson 2000: Chapters 2,3). However, rational beliefs with no factual defeaters are even more robust to new evidence, but knowledge is still more interesting.

¹² See fn. 12.

have determined which conditions must obtain for the subject's reason to be good in both senses, we still need to determine what it is for her reason to be good in both senses. Each species of the genus gives us an account of which conditions must obtain for the subject to have a subjective/objective normative reason. They do this by giving an account of how the subject's justifying reasons need to stack up against a body of potentially countervailing reasons to be sufficient. In the case of subjective normative reasons, the body of potentially countervailing reasons is the subject's total evidence. In the case of objective normative reasons, the body of potentially countervailing reasons is the totality of the facts.

However, this doesn't tell us what it is to have a good reason in either sense. To see the point, consider the following example. Suppose you can determine whether you have high cholesterol by spitting on a stick and putting it in chemical solution A and you can determine if you have lupus by spitting on that same stick and placing it in chemical solution B. Does this show that there is an interesting connection between what it is to have lupus and what it is to have high cholesterol? No.

Similarly, it may be the case that you can determine whether a subject has a subjective normative reason by performing a procedure relating her reason to one body of evidence and you can determine if she has an objective normative reason by performing the same procedure relative to the totality of the facts. But this doesn't show that subjective normative reasons and objective normative reasons have any more explanatory connection than lupus and high cholesterol would have if chemical solutions A and B were actual. Because he recognizes this point, Schroeder offers the following: a subjective normative reason is a reason that would be objective normative if the subject's beliefs were true (2007, p. 14, 2008, 2015).

Some counterfactual in the neighborhood is true. I prefer the following: if the world were to be as the subject justifiably takes it to be, then her subjective normative reasons would be objective normative reasons.¹³ Either way, the counterfactual doesn't explain what needs to be explained. As an account of the conditions that must obtain for reasons to be subjective normative, this is fine. However, it doesn't tell us what it is to be a subjective normative reason. It doesn't give an adequate account of their conceptual connection with objective normative reasons. As with Audi's proposal, the problem stems from the fact that the analysis is stated as a counterfactual (though this time not for the reasons identified by Shope (1983)). If the subject's beliefs were true, the subjective normative reason would be an objective normative reason (*mutatis mutandis* for my proposed revision). This tells us nothing about the standing of her reasons in the actual world. It tells us nothing about how subjective normative reasons are connected to objective normative reasons in the actual world. If wishes were horses, then beggars would ride. Wishes aren't horses. So, what is the connection between wishes and horses in the actual world? There isn't one! If the subject's beliefs were true, then her subjective normative reasons would be objective normative reasons. What about when they aren't true? What is the connection between subjective and objective normative reasons then? Schroeder's account doesn't help us here and neither

¹³ Putting it this way runs afoul of the reasons-first order of explanation. There is a more cumbersome way of making the same point that does not: if the world were as the subject rationally takes it to be (where rationality is determined by the balance of subjective normative reasons), then her subjective normative reason would be an objective normative reason.

does any other counterfactual account. Perhaps this is why in more recent work he remains agnostic about the connection between the two (Schroeder, 2021, p. 45).¹⁴ But there surely is still a connection between subjective normative reasons and objective normative reasons in the actual world, even when the subject's beliefs are false, and the connection isn't identity. Surely there is a connection *from the subject's point of view*.

This is what explains why Schroeder's counterfactual is true. If the subject's beliefs were true, her subjective normative reason would be an objective normative reason. This is because the subject's beliefs are the facts, so far as she can tell. Similarly, her justified beliefs (i.e., those for which she has sufficient subjective normative reasons) are the facts, so far as someone with her total evidence can tell.¹⁵ So, her subjective normative reasons for thinking that those are the facts are just objective normative reasons, so far as one with her total evidence can tell. If the facts are as she takes them to be, then her subjective normative reasons are not misleading. In that case, they are objective normative reasons. But even if they aren't in fact objective normative reasons because her total evidence is misleading, it remains the case that her subjective normative reasons appear to be objective normative reasons to one with her total evidence. That is what it is to be a subjective normative reason and it is what explains why the counterfactual is true. In general, the counterfactual is explained in terms of the actual. This is no exception.

Of course, more needs to be said about the sense of "appear" at play here, and more will be said in the next section. The key point is that we don't have a non-gerrymandered account of knowledge until we can articulate the connection between what it is to be a subjective normative reason and what it is to be an objective normative reason. None of the species mentioned above do that. They give us an account of which conditions must obtain for reasons to be normative in both senses. But they don't give an account of what it is to be normative in either sense that makes clear the connection between the two senses of "normative reason" because they don't tell us what it is to be a subjective normative reason in a way that explains the connection.

4 Good reasons are apparent to the knowing subject

I will begin by offering some further motivation for the claim that what it is to be a subjective normative reason is to be an apparent objective normative reason to one in the subject's epistemic position (i.e., with her total evidence). It seems clear that there is an important sense in which the subject in a Gettier case has good reasons for her Gettiered belief (the same goes for justified false beliefs). Reasons, by their very nature, rationalize belief. This makes their role in rationalizing belief essential to them. When reasons rationalize belief, they do something they are supposed to do *qua* reason. That makes them good reasons, at least in one important sense. However, this is not all that reasons are supposed to do. In addition to standing up to scrutiny in light of the subject's possessed evidence, they are also supposed to stand up to

¹⁴ Though see (Schroeder 2021, p. 69) for a possible lapse into his earlier view.

¹⁵ Strictly speaking, these are her justifiable beliefs. I ignore the subtlety here to streamline exposition.

scrutiny in light of the facts. When reasons don't stand up to scrutiny in light of the facts, they fail to do something they are supposed to do *qua* reasons. For this reason, it would make perfect sense for an enlightened third-party to say the following about a Gettier subject, "Her reasons are no good, although she is in no position to tell." The third-party might qualify the remark by noting that the Gettier subject's reason was good *for her* at the time. That is, it was not good in the objective normative sense but perfectly fine in the subjective normative sense.

These remarks indicate not only that there are two senses of "good reason" or "normative reason", but also gesture in the direction of how they relate. In the last section, I said that what it is to be a subjective normative reason is to be apparently good in the objective normative sense. As we just saw, it would be quite natural to say that a Gettier subject's reasons weren't any good, although she was in no position to tell. That is, she didn't have an objective normative reason, although she reasonably thought otherwise. I say this is why she had a subjective normative reason. To have a subjective normative reason just is to apparently have an objective normative reason. This explains why the remark from the enlightened third party makes sense.

I don't mean that for a reason to be good in the subjective normative sense the subject must also have a higher-order reason to think that it is good in the objective normative sense.¹⁶ If this seems puzzling, the following remark may help. The question of what to believe is transparent to the question of what is the case (cf. Evans, 1982; Moran, 2001). This does not mean that it is impossible to believe what is not the case. It just means that when you are trying to figure out what to believe, you do so by trying to figure out what is the case. I urge that a similar point applies to figuring out what your reasons are. When you are trying to figure out if your perceptual evidence is really a reason for you to form a belief about the address of the house you just drove past, you try to determine whether there is in fact some reason to doubt your senses in this situation, i.e., an objective normative reason. You don't first figure out the subjective normative reasons and then try to build a bridge to the objectively good reasons through meta-reasons. Rather, you just figure out what the objective normative reasons are. The fact that this takes place within the inquirer's limited perspective does not mean that they first take stock of how things are from within that perspective and then try to see how it lines up with the facts (*pace* Fumerton, 1995, p. 174; McClelland & Chihara, 1985, p. 81).

In fact, Fumerton, BonJour, and Pollock & Cruz's remarks from earlier have things backwards. They think that from the first-person perspective you check to see if your reasons are good in the subjective normative sense. Then there is no further inquiry into how they line up with the facts. You can't jettison your perspective to go check. So, they think the theory of good reasons in the subjective normative sense (i.e., justification as they understand it) is the heart of epistemology. How reasons line up with the facts is just something you have to tip your hat to when knowledge comes up.

I urge that what you care about from the first-person perspective is how your reasons line up with the facts, i.e., goodness in the objective normative sense (cf. Littlejohn, 2012). Once you've determined that as well as someone in your position can, there is no further question about your subjective normative reasons. The goodness of reasons in

¹⁶ Thanks to an anonymous referee for making clear the importance of driving this point home.

the objective normative sense is the heart of epistemology. Goodness in the subjective normative sense is just a finite subject's best estimate of that. There is nothing left over to which we must reluctantly tip our hats.

A good reason in the subjective normative sense is apparently (to one in the subject's position) a good reason in the objective normative sense. To see this, try to consider a subject for whom *R* is a subjective normative reason, i.e., it rationalizes some belief for her, but for whom it is also reasonable to believe that, if she were better informed, *R* would no longer be able to rationalize that belief for her. You can't conceive of such a subject. If the second conjunct is true, the first isn't. If it is reasonable for her to think that *R* only rationalizes her beliefs because she is in a benighted position, then *R* isn't really a good reason for belief to one in her position.¹⁷ So, it doesn't really rationalize her beliefs. It is a merely *prima facie* reason defeated by higher-order evidence.

Conversely, the more reasonable it is to think that your reasons will continue to support your belief relative to more comprehensive bodies of evidence, the weightier they become *qua* subjective normative reasons. The professor in Lehrer's Gettier case was reasonable in taking the student's testimony to be a good reason in the objective normative sense. Rarely does testimony about car ownership fold under further inquiry. However, in this case it did. So, it was not a good reason in the objective normative sense, but it was a good reason in the subjective normative sense. This is because it was reasonable for the subject to take it as good in the objective normative sense. If she had even more reason to think that the testimonial reason would survive further inquiry (e.g., misleading evidence that the student is more honest than average), then her reasons would be even better subjective normative reasons.

The key takeaway is that my account does not make subjective normative reasons harder to come by due to their identification with apparent objective normative reasons. I am not saying that having a subjective normative reason requires the absence of countervailing reasons in the subject's possessed evidence and, *additionally*, reasons to think there aren't countervailing reasons out there of which she is presently unaware. Rather, by identifying subjective normative reasons with apparent objective normative reasons, I am saying that the absence of countervailing reasons in the subject's possessed evidence *just is* apparent lack of sufficiently weighty countervailing reasons "out there". This is the upshot of the transparency point. The subject's beliefs are her take on the facts and her justifying reasons are her take on the reasons that will ultimately stand up to scrutiny in light of the facts.

This explains why "Custer died at Little Big Horn, but I don't know that he did" sounds bad (Littlejohn, 2012, p. 2; cf. Smithies, 2012). You can't rationally believe the conjunction. If you rationally believe the first conjunct, then you are committed to the claim that your reasons are of the sort that turn rational true beliefs into knowledge, i.e., objective normative reasons. You are committed to this because you take yourself to have a subjective normative reason (otherwise asserting the first conjunct would be bad), and that is just an apparent objective normative reason. But if your apparent objective normative reason really is as it appears to be, then you know the first conjunct.

This does not generate a regress. When the subject uses *r* as a reason for believing *p* by basing belief that *p* on *r*, she implicitly regards *r* as a good (objective normative)

¹⁷ Ballantyne (2015) makes much of this point.

reason for believing p . This is the transparency point, with the extra note that one's attitude toward one's reasons can be implicit through basing. She does so reasonably if, so far as one can tell from her perspective, r is a good reason for p . When this condition obtains, it will sometimes be because she has some further reason for so regarding r . In other cases, r will be a *prima facie* reason for believing p . That is, it will not depend positively on other reasons for its own status as a good reason. Rather, it depends negatively on them: It doesn't depend on them for support, only the absence of defeat. This halts the regress. I won't argue further for the legitimacy of the distinction between positive and negative epistemic dependence (see Audi, 1993 for an argument).

The distinction between positive and negative epistemic dependence also applies to your attitudes towards your reasons. Your belief that you have hands when you see them in ostensibly normal conditions depends negatively on the absence of reasons to think your perceptual states are misleading in this instance. It does not depend positively on independent reasons to think they are not illusory. Similarly, it is reasonable to take your perceptual reason to be an objective normative reason in the absence of reasons to think otherwise. If it is not reasonable to take your reason to be an objective normative reason, then you cannot reasonably believe on its basis. If you are not in a position to reasonably believe that your reasons will withstand scrutiny in light of the facts, then they aren't even good reasons relative to your total evidence (i.e., they are not even subjective normative reasons). Why is it reasonable for the Professor in Lehrer's example to believe that the student owns a Ferrari? Because he said so, and it is reasonable for her to believe that her testimonial reason is not trumped or mitigated by further facts. Testimony of this sort usually holds up. If it weren't reasonable to think the reason will hold up in the face of the facts, then she would have no business believing on its basis in the first place. She should suspend judgment instead.

The problem with which we began is that the connection between (1) and (2) is unclear. Without a clear account of the connection, the concept KNOWLEDGE appears gerrymandered. We have now seen how to avoid this result. Reasons that do (1) are (by their very nature) reasons that apparently do (2). We can now say a bit more to alleviate the gerrymandering worry. In the last section we saw a family of approaches according to which (1) is a matter of how the subject's reasons square with her possessed evidence and (2) is a matter of how they square with the facts. This helps to the extent that it shows (1) and (2) to be complementary perspectives on a single thing: the subject's reasons. However, the gerrymandering worry still rears its head. Unless more is said, it appears that when we assess the subject's reasons relative to her possessed evidence, we are checking for one property: rationality. When we assess her reasons relative to the facts, it isn't immediately clear that we are looking at that property from a different perspective as opposed to looking at an entirely different property. Schroeder, for instance, calls the dimension of assessment relevant to (2) "alethic" (2015). This gives the impression that we are using one procedure to check for two very different things: one alethic and the other properly epistemic. This is the point made earlier by the lupus/cholesterol example. We saw that Schroeder's counterfactual analysis of subjective reasons in terms of objective reasons doesn't help matters.

Notice, however, that in assessing the subject's reasons relative to the totality of the facts, we are treating the totality of the facts like a body of evidence, like that possessed by any finite subject at any given time. The body of evidence composed of all and only the facts is the (regulatory) ideal end point of inquiry. It is the point at which there is nothing further into which one can inquire. If the subject's reasons square with the totality of the facts (according to the procedure laid down by whichever species of the genus considered in the last section proves to be correct) then her reasons are good relative to the perspective that is the ideal end point of inquiry. This enables us to better see the explanatory connection between (1) and (2). Assessing the subject's reasons relative to her own perspective (i.e., total evidence) and the perspective composed of all and only the facts is not like dipping her beliefs in unrelated chemical solutions. Rather, it is to assess the subject's reasons relative to her own perspective and also relative to the ideal end point of the activity in which she is engaged as she reasons.

The view offered here can, with a slight emendation, also address another problem for Schroeder's view (indeed, other species of the genus as well). It is possible for reasons to be both subjectively and objectively sufficient, and a belief to be based on them, but that belief might still fail to be knowledge. I diagnose the problem in the following way: it is possible for good reasons to be apparent to a subject without the goodness of those reasons being apparent to her. The latter is necessary for knowledge.

This point is illuminated by Turri's (2010) example in which the subject justifiably believes that p , and that if p , then q . Those are the reasons for which she believes q . These are good reasons (in both the subjective and objective normative senses) to believe that q . However, she doesn't know that q . Instead of inferring q from those reasons using *modus ponens*, she infers it using the inference rule: infer q from any two premises. The reasons themselves are perfectly fine, the subject's use of those reasons is problematic (Ibid).

On my account, good reasons are apparent to this subject, but the goodness of those reasons is not. What makes those reasons good is the fact that they are fit to figure in a *modus ponens* argument. One needn't form an explicit belief to suitably appreciate this fact. One suitably appreciates this fact just by basing one's belief that q on p and $p \rightarrow q$ by inferring q from these premises in accordance with *modus ponens*. By inferring one's belief that q in this way, one suitably appreciates the goodness of q by appreciating what it is about one's reasons that makes them good reasons to believe q : their role in *modus ponens*. That is, it is not just apparent that these reasons are good: what makes them good is apparent as well. So, strictly speaking, the goodness of good reasons is apparent to the knowing subject. That is a mouthful, so I will stick with the punchier "[good reasons are apparent to the knowing subject](#)" as the mantra of the paper. However, the official position is nonetheless that their goodness must also be apparent.

We can now state what our interest in knowledge is all about. We are interested in whether good reasons are apparent to the subject. We are interested in whether her reasons really are as they appear to her. Things really being as they appear to one is clearly valuable (cf. Williamson, 2000, p 40). (2) is a matter of whether the subject's reasons really are good and (1) is a matter of whether they are apparently good. They are both about the goodness of reasons. Interest in their apparent goodness just is an

interest in their goodness, from the point of view of a finite and possibly mistaken perspective.

I should also mention that the subject having good reasons in the objective normative sense is truth-entailing, so there is no need to mention it separately.¹⁸ The account I recommend is still fallibilist. The fallibilist maintains that we can know on the basis of defeasible reasons. I agree. The knower's reasons are defeasible, they are just not, in fact, defeated (in the same way my books are flammable but not, in fact, on fire).

When you believe something, you know it just in case your belief is based on good reasons apparent to you. This sounds so obvious that a non-philosopher would think it is hardly worth stating. This is a point in favor of the view. It shows us how to say what we were inclined to say before we entered the epistemology classroom in a way that is consistent with the lessons epistemology has taught us.

Furthermore, the metaepistemological view that emerges doesn't require us to say that epistemology is really about one aspect of knowledge rather than the other. Epistemology is about reasons to believe (i.e., our metaepistemology is "reasons-first"). Epistemic assessment is about the goodness of those reasons. Once we see what knowledge is, we see that it is the unified, interesting phenomenon we implicitly take it to be. We have made sense of knowledge in terms of a kind of assessment that is an important part of how we understand ourselves. The metaepistemological view I recommend enables us to see epistemology as being unified by a single kind of assessment, one everyone implicitly understands before entering the epistemology classroom.

5 Conclusion

I have argued that the role reasons play in rationalizing belief and the role they play in turning true belief into knowledge can only be understood in tandem. The picture that emerged is that a subject knows just in case her belief is based on good reasons apparent to her. This enables us to make sense of and reflectively endorse our practice of traffick-ing in knowledge attributions. What appears to us as a unified and significant aspect of our lives is what it appears to be. Certain tendencies in epistemology make it difficult to see this. Those tendencies are difficult to avoid. Once we see that the components of knowledge can come apart, it is difficult to see what unites them. Recognition of this difficulty, even if just inchoate, can incline us to think that the Gettier intuition is faulty (Olsson, 2015; Weatherson, 2003; Weinberg, 2017), or that knowledge shouldn't be analyzed in terms of rationality plus something else (e.g., Zagzebski, 1994, 2017), or perhaps that we shouldn't try to analyze it at all (Williamson, 2000). I hope to have shown that we can understand knowledge in terms of a prior understanding of reasons without our analysis turning into the "ad hoc sprawl" (Williamson, 2000, p. 31) that gives these alternatives much of their appeal. There is nothing ad hoc about the phenomenon of one's reasons really being as they appear.

Acknowledgements I'd like to thank Beth Barker, Regina Hurley, Jennifer Lackey, Nate Laufer, Kathryn Pogin, Baron Reed, Tamler Sommers, Justin Zacek and two (incredibly helpful) anonymous referees for comments and helpful conversations. Most of all, thanks to Sandy Goldberg, for reading what must have

¹⁸ Rosenberg (2002) has a proposal that is similar in this respect, although relativist in a way mine isn't.

been a dozen drafts over several years, and (somehow) getting contagiously excited each time he read a new one.

Funding I have received no external funding for this work.

Declarations

Conflict of interest I have no conflicting interests to declare.

References

- Audi, R. (1993). *The structure of justification*. Cambridge University Press.
- Ballantyne, N. (2015). The significance of unpossessed evidence. *The Philosophical Quarterly*, 65(260), 315–335.
- Barker, J. (1976). What you don't know won't hurt you? *American Philosophical Quarterly*, 13, 303–308.
- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- BonJour, L. (2003). In L. BonJour, & E. Sosa (Eds.), *Epistemic justification*. Blackwell Publishing.
- BonJour, L. (2010). *Epistemology: Classic problems and contemporary responses* (2nd ed.). Rowman & Littlefield Publishing.
- Boyle, M. (2012). Essentially rational animals. In G. Abel & J. Conant (Eds.), *Rethinking epistemology*. Walter de Gruyter Verlag. DeGruyter.
- de Almeida, C., & Fett, J. R. (2016). Defeasibility and Gettierization: A reminder. *Australasian Journal of Philosophy*, 94(1), 152–169.
- Evans, G. (1982). In J. McDowell (Ed.), *The varieties of reference*. Clarendon Press.
- Fumerton, R. (1995). *Metaepistemology and skepticism*. Rowman & Littlefield Press.
- Gettier, E. (1963). Is justified, true belief knowledge? *Analysis*, 23(6), 121–123.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Hilpinen, R. (1971). Knowledge & justification. *Ajatus*, 33(7), 39.
- Horty, J. (2007). Reasons as defaults. *Philosophers' Imprint*, 7, 1–28.
- Horty, J. (2012). *Reasons as defaults*. Oxford University Press.
- Johnsen, B. (1974). Knowledge. *Philosophical Studies*, 25(4), 273–282.
- Kaplan, M. (1985). It's not what you know that counts. *Journal of Philosophy*, 82(7), 350–363.
- Kirkham, R. (1984). Does the Gettier problem rest on a mistake? *Mind*, 93(372), 501–513.
- Klein, P. (1971). A proposed definition of propositional knowledge. *Journal of Philosophy*, 68(16), 471–482.
- Klein, P. (1976). Knowledge, causality, indefeasibility. *Journal of Philosophy*, 73(20), 792–812.
- Klein, P. (1980). Misleading evidence and the restoration of justification. *Philosophical Studies*, 37(1), 81–89.
- Klein, P. (2017). The nature of knowledge. In R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining knowledge: New essays on the Gettier problem* (pp. 35–57). Oxford University Press.
- Kornblith, H. (2011). *On reflection*. Oxford University Press.
- Korsgard, C. (1996). *The sources of normativity*. Harvard University Press.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Kvanvig, J. L. (2017). Lessons from Gettier. In R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining knowledge: New essays on the Gettier problem* (pp. 152–169). Oxford University Press.
- Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25(5), 168–175.
- Lehrer, K. (1970). The fourth condition for knowledge: A defense. *Review of Metaphysics*, 24(1), 122–128.
- Lehrer, K., & Paxson, T. (1968). Knowledge: Undefeated justified true belief. *Journal of Philosophy*, 66(8), 225–237.
- Littlejohn, C. (2012). *Justification and the truth-connection*. Cambridge University Press.
- McClelland, J., & Chihara, C. (1985). The surprise examination paradox. *Journal of Philosophical Logic*, 4(1), 71–89.

- Moran, R. (2001). *Authority & estrangement: An essay on self-knowledge*. Princeton University Press.
- Moser, P. K. (1988). Review of John Pollock's *Contemporary Theories of Knowledge*. *British Journal for the Philosophy of Science*, 39(1), 131–140.
- Moser, P. K. (1989). *Knowledge & evidence*. Cambridge 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(1), 57–72.
- Pollock, J. (1986). *Contemporary theories of knowledge* (1st ed.). Rowman & Littlefield.
- Pollock, J., & Cruz, J. (1999). *Contemporary theories of knowledge* (2nd ed.). Rowman & Littlefield.
- Putnam, H. (1975). The meaning of 'meaning.' *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
- Rosenberg, J. (2002). *Thinking about knowing*. Oxford University Press.
- Scanlon, T. M. (1998). *What we owe to each other*. The Belknap Press of Harvard University Press.
- Schroeder, M. (2015). Knowledge is belief for sufficient (objective and subjective) reason. In *Oxford studies in epistemology* (Vol. 5, pp. 226–252).
- Schroeder, M. (2007). *Slaves of the passions*. Oxford University Press.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57–71.
- Schroeder, M. (2021). *Reasons first*. Oxford University Press.
- Sellars, W. (1956/1997). *Empiricism & the philosophy of mind*. Harvard University Press.
- Shope, R. (1983). *The analysis of knowing: A decade of research*. Princeton University Press.
- Smithies, D. (2012). Moore's paradox and the accessibility of justification. *Philosophy and Phenomenological Research*, 85(2), 273–300.
- Smithies, D. (2015). Why justification matters. In D. Henderson & J. Greco (Eds.), *Epistemic evaluation: Purposeful epistemology*. Oxford University Press.
- Smithies, D. (2016). Reflection on: On reflection. *Analysis*, 76(1), 55–69.
- Steup, M. (1996). *An introduction to contemporary epistemology*. Prentice Hall.
- Swain, M. (1974). Epistemic defeasibility. *American Philosophical Quarterly*, 11(1), 15–25.
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80(2), 312–326.
- Weatherson, B. (2003). What good are counterexamples? *Philosophical Studies*, 115(1), 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.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *Philosophical Quarterly*, 44(174), 65–73.
- Zagzebski, L. (2017). The Lesson of Gettier. In R. Borges, C. de Almeida, & P. Klein (Eds.), *Explaining knowledge: New essays on the Gettier problem* (pp. 179–191). Oxford University Press.

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.