Topoi (2023) 42:83–89
https://doi.org/10.1007/s11245-022-09845-7

Sensitivity, Safety, and Brains in Vats

Haicheng Zhao[1]

Accepted: 2 November 2022 / Published online: 11 January 2023

Abstract
Both sensitivity and safety theorists concur that their accounts should be relativized to the same method that one employs in the actual world. However, properly individuating methods has proven to be a tricky matter. In this regard, Nozick (Philosophical Explanations, Harvard University Press, Cambridge, 1981) proposes a Same-Experience-Same-Method Principle: if the experiences associated with two method tokens are the same, they are of the same type of method. This principle, however, has been widely rejected by recent safety and sensitivity theorists. In this paper, I raise an argument in favor of Nozick: not endorsing the principle leads to some rather implausible consequences when certain skeptical scenarios are considered—i.e., scenarios in which skeptical possibilities are 'close'. Additionally, this argument reveals some important lessons about skepticism in general and the place of modal accounts of knowledge in the internalism/externalism debate.

Keywords Sensitivity · Safety · Skepticism · Method

1 Introduction

Many concur that knowledge excludes luck. A popular way to accommodate this intuition is to adopt a modal condition on knowledge. The general idea here is that if one knows that *p*, one does not merely form a true belief that *p* in the actual world, but the belief continues to 'track' truths in a proper range of possibilities. Two prominent modal accounts are sensitivity and safety, which can roughly be formulated as follows:

Safety: If S knows that *p*, then S's belief that *p* is true in nearly all close possible worlds in which S continues to believe that *p*.

Sensitivity: If S knows that *p*, then in the closest possible worlds in which *p* is false, S does not believe that *p*.[1]

As well known, such accounts must add a 'method' clause. One often-cited example in this regard is Nozick's (1981, p. 181) grandmother case. Suppose that a grandmother sees clearly that her grandson is well. She thus believes that he is well. However, if the grandson were sick, the relatives would still tell her that he is well, in order to spare her upset.

Both *Safety* and *Sensitivity* yield the wrong verdict here. Suppose that the grandson could easily have gotten sick. This does not seem to prevent the grandmother from knowing that he is well upon seeing him so. However, according to *Safety*, she does not know. Also, if the grandson were sick, the grandmother would still believe that he is well, via her relatives' testimony. So *Sensitivity* also predicts that she does not know.

Consequently, many prefer to add a method clause to *Safety* and *Sensitivity*:

Method Safety: If S knows that *p* via method m of type T, then S's belief that *p* is true

✉ Haicheng Zhao
haichengzhao90@gmail.com

[1] Department of Philosophy, Xiamen University, 422 Siming S Rd, Siming District, Xiamen 361005, Fujian, China

[1] Strictly speaking, *Safety* and *Sensitivity* (as well as *Method Safety* and *Method Sensitivity* below) are not the definitions of the safety principle and the sensitivity principle in the literature. They express the theses that these principles are necessary conditions for knowledge. In addition, safety and sensitivity principles can be formulated in different ways. For alternative versions of safety, see e.g., Beddor and Pavase (2020), Broncano-Berrocal (2018), Hirvelä (2019), Peet and Pitcovski (2018), Pritchard (2007, p. 292, 2012, pp. 256–257), Smith (2016). For alternative versions of sensitivity, see e.g., Becker (2007), Broncano-Berrocal (2018), Cross (2010), DeRose (2010), Roush (2007).

in nearly all close possible worlds in which S continues to believe that *p* via method m* that belongs to type T.

Method Sensitivity: If S knows that *p* via method m of type T, then in the closest possible worlds in which *p* is false and S uses method m* that belongs to T to determine whether or not *p*, S does not believe that *p* via m*.

By restricting the relevant possibilities to those where the same type of method is being employed, these accounts nicely address the above counterexample. Starting with *Method Safety*, even if the grandmother forms a false belief that her grandson is well in some close worlds, she does not form the belief via the same perceptual method. Rather, in those worlds the belief is formed via testimony. Per *Method Safety*, these worlds are irrelevant. Therefore, her belief is safe, as desired. According to *Method Sensitivity*, in closest worlds in which her grandson is not well and she uses her *perceptual method* to determine if he is well, she would not believe that he is well. Thus, *Method Sensitivity* leads to the result that she knows as well.[2]

2 Nozick's SESM and Its Discontents

However, invoking methods immediately raises the tricky question about how to individuate a method. Different individuations of methods may well entail that different sets of possibilities are being picked out, thus resulting in extensionally different versions of *Method Safety* and *Method Sensitivity*. Consequently, the explanatory power of these accounts substantially depends on a proper account of methods. In this regard, Nozick (1981, pp. 184–185) proposes to individuate methods in an internal fashion:

> Usually, a method will have a final upshot in experience on which the belief is based, such as visual experience, and then (a) no method without this upshot is the same method, and (b) any method experientially the same, the same "from the inside", will count as the same method.

With this approach to methods, Nozick concludes that we and brains in vats share the same method (p. 185). More generally, the idea can be encapsulated as follows:

Same Experience Same Method (SESM): For any two method tokens m1 and m2, if the experiences associated with these method tokens are the same, they are of the same type of method.

Before proceeding, one clarification is in order. Note that *SESM* is a principle for determining when two method *tokens* are of the same type (where 'token' can be considered as belief-formation *event*) (Cf. Alfano 2009; Broncano-Berrocal 2014; Hirvelä 2019). Crucially, it is *not* a principle for deciding when two method types are the same. To illustrate the differences here, consider two method types: vision (M1) and taking a hallucinatory drug (M2). Clearly, M1 and M2 are not the same type of method. But this is compatible with arguing that some method *tokens* generated by M1 and M2 belong to the same type. For instance, suppose that one has a vivid zebra-like-experience and believes that there is a zebra by M1, whereas another subject has the same experience (though hallucinatory) and forms the same belief by M2. Then, according to *SESM*, the method tokens involved in these two cases are of the same type.[3]

Now, *SESM* is essential to Nozick's sensitivity-based reply to skepticism. According to him, I do not know that I am not a BIV. This is made possible by *SESM*. As I am an experiential duplicate of my BIV counterpart, both the BIV and I share the same method according to *SESM*. So, if I were a BIV and appealed to *the same (experiential) method* to determine if I am a BIV, I would still believe that I am not a BIV. *Method Safety*, on the other hand, would generate the Moorean conclusion that I know I am not a BIV, even if one accepts *SESM*: Assuming that our world is just the way it is, without any BIVs or the like, the BIV worlds are quite dissimilar compared to the actual world. So, there are no close worlds in which I appeal to the same experiential method but end up believing falsely that I am not a BIV.

Recently, however, philosophers from both the sensitivity and the safety camps have rejected *SESM*.[4] Williamson (2000, 2009), a safety theorist, prefers to individuate methods externally. He says: "I had in mind a very liberal conception, on which the basis of a belief includes the specific causal process leading to it and the relevant causal background." (2009, p. 307) Furthermore, commenting on Goldman's (1976) dachshund-wolf case, he claims that seeing a

[2] Incidentally, in some cases one may believe that *p* via multiple methods, and sometimes one of these methods may outweigh others. For discussions of these issues, see Nozick (1981, pp. 179–185).

[3] Relatedly, *SESM* does not claim that two method types must deliver the same experiences every time upon the employment of these method types, in order for them to be counted as the same type. I thank an anonymous reviewer for suggesting me to clarify these issues.

[4] As a terminological note, instead of using the term 'method', Pritchard prefers to use the term 'way of belief formation' and Williamson prefers to use the term 'basis'. I take these differences to be merely verbal. What is essential is how one individuates method (basis/way of belief formation).

dachshund and seeing a wolf involve large enough external differences, so that the methods of seeing them are different (2009, p. 307). If so, it follows that on Williamson's view, we and BIVs are not using the same method. Presumably, the causal processes leading to our beliefs and BIVs' beliefs are no less different than the processes of seeing a wolf and a dachshund.[5]

Pritchard, another safety theorist, also abandons *SESM*. He claims that "the 'way' in which the belief is actually formed needs to be individuated *externally* rather than *internally*…" (2005, p. 152). Pritchard does not say much about what such an external approach to methods amounts to, but commenting on BIV scenarios, he claims that "…the way in which the belief was actually formed is, it seems, unavailable to the agent in these [BIV] worlds" (2005, p. 160; see also his 2008). Apparently, this takes into account the external differences between our methods and a BIV's. *Ex hypothesi*, BIVs form their beliefs by the scientists feeding their experiences, whereas our beliefs are typically produced by veridical experiences.

Indeed, most recent safety theorists do individuate methods in an externalist fashion (See also Broncano-Berrocal 2014; Grundmann 2020; Hirvelä 2019).[6] And regardless of the differences in details, one commonality among these theorists is that their accounts of methods entail the rejection of *SESM*.

In the sensitivity camp, Black (2002, 2008) also rejects *SESM*. He thinks that methods should be characterized at least partly 'from the outside'. Thus, in perceptual cases, non-experiential physical items, such as retinas, tympanic membranes, are considered as part of a method. Based on this, Black gives a sensitivity-based neo-Moorean reply to the sceptic: I know that I am not a BIV because my actual method of believing so—which Black takes to be the result of perceptual processes—is not available to the BIV. By stipulation, BIVs do not have retinas, tympanic membranes, etc., so there are no worlds in which I am a BIV and I appeal to the same perceptual method to determine if I am one. *Method Sensitivity* is trivially satisfied.[7]

In what follows, I will give an argument in support of *SESM*. As will be shown, abandoning *SESM* and adopting externalist approaches to methods will lead to rather implausible forms of *Method Safety* and *Method Sensitivity*.

3 Close Skeptical Possibilities

Consider a possible world w_1. In w_1 BIVs are actual. A large group of scientists have secretly invented all the equipment needed for BIVs. Conspiring with a criminal organization whose daily task is to kidnap as many people as possible, the scientists have actively engaged in their BIV-operations. Many have already become victims. However, since the criminals have covered up their crimes perfectly, the police have no clue about the missing people. All the victims are turned into BIVs by the scientists—a procedure that guarantees that all of their memories regarding kidnapping are erased and they cannot subjectively tell anything abnormal about their situations. From their own perspectives, things just continue the way they are.

Suzy, an inhabitant of w_1, is a college student. She has never heard of the actual BIV operations nor the criminal organization. Today, she learns in a philosophy class that BIV is at least a possibility. But, just as many people do, Suzy believes that she is not a BIV, although she could very easily have been kidnapped and become one. Indeed, the criminal organization has recently become very interested in kidnapping college students in Suzy's area. At the moment Suzy forms the belief that she is not a BIV, many of her fellow students are kidnapped.

Intuitively, Suzy does not know that she is not a BIV. After all, given that she could easily have become a BIV, it is just a matter of luck that she acquires a true belief that she is not one (More on this later). Indeed, many concede that even *we* do not know we are not BIVs. For these philosophers, it should appear quite plausible that Suzy does not know either. Moreover, there are philosophers who think that we *know* we are not BIVs. Safety theorists (Sosa 1999, 2007; Pritchard 2005), as 'neo-Mooreans', claim that insofar as there are no BIVs in our world, the BIV world is relatively remote, so that the possibility of BIV does not affect the

[5] It is worth noting that Williamson, as an advocate of 'knowledge-first epistemology', does not advance safety as a non-circular necessary condition for knowledge. As such, he thinks that we may use our intuition about knowledge to make judgments about whether one's method in a case is similar enough to the actual method (Williamson 2009). My main target in this paper is the majority of safety theorists who advance safety as a non-circular necessary condition for knowledge and who also reject *SESM*.

[6] Grundmann (2020, p. 5173) claims that "…the method is externally individuated if it does not supervene on the agent's reactive dispositions but is partly constituted by the actual nature of the external conditions." This naturally implies the rejection of *SESM*: two experientially same method tokens could belong to different types if, on Grundmann's account, they involve sufficiently large differences in external conditions. Furthermore, both Broncano-Berrocal (2014) and Hirvelä (2019) claim that two method tokens are of the same type just in case they are globally reliable to the same degree with respect to the same field of propositions and the same range of circumstances. This entails the rejection of *SESM* as well: Our beliefs and BIVs' beliefs are formed via the same experiences but they differ radically with regard to reliability.

[7] For more discussions on how methods should be individuated for sensitivity or safety accounts, see Alfano (2009), Becker (2008, 2012), Black and Murphy (2007), Bogardus and Marxen (2014), Greco (2016).

safety of our beliefs that we are not BIVs. But even if we assume that this reasoning is cogent, these theorists should still concur that Suzy does not know. For in her situation, the BIV-possibility is precisely a very close one. Thus, if safety theorists insist that Suzy knows, this just undermines their argument to the effect that *we* know we are not BIVs because such possibilities are remote. Consistency requires that if we know that we are not BIVs because such possibilities are remote, then Suzy fails to know because such possibilities are rather close.

Additionally, given Suzy's situation, her ignorance is unaffected by the possible methods of her belief-formation. There is no agreement about how one believes the proposition that one is not a BIV. Some take it that one may believe so based on perceptions or inferences from perceptions (Black 2002; Pritchard 2008). Some think that such a proposition can be believed (and known) based on what we take to be technically possible in our world (Cf. Zalabardo 2017). Still some claim that one may believe so based on one's epistemic ability—an ability which recognizes that BIV possibilities could not easily obtain (Kelp 2021). Potentially, these are all possible ways of forming a not-BIV belief. But no matter which of these methods Suzy employs, it seems clear that the luck involved in her belief prevents her from knowing: were things to be just slightly different, she would have been kidnapped and formed a false belief that she is not a BIV. The luck involved here is precisely the kind of knowledge-precluding element that modal epistemologists seek to eliminate.

However, with the above external approaches to methods—approaches that reject *SESM*—neither *Method Safety* nor *Method Sensitivity* can explain why Suzy does not know. Consider *Method Safety* first, although for Suzy there are rather close worlds in which she becomes a BIV and forms a false belief that she is not a BIV, these are worlds in which her method involves rather significant external differences. Thus, on the external approach to methods, these close worlds are irrelevant for epistemic evaluation. Consequently, proponents of *Method Safety* have to admit that Suzy's belief is safe, contrary to the intuition that she does not know. This result can be avoided by adopting *SESM* and conceding that Suzy's method is of the same type as her BIV-counterpart's.[8]

Likewise, with *Method Sensitivity* and external approaches to methods, one has to acknowledge that Suzy knows. This is because there are no closest worlds in which Suzy is a BIV *and* she appeals to the same external method. This conjunction is a counterpossible, which implies that Suzy's belief is trivially sensitive. But with *SESM*, the closest antecedent worlds would be those in which Suzy is a BIV and she appeals to the same experiential method as in the actual world to determine if she is a BIV. These are precisely worlds in which she is kidnapped by the criminals and ends up being a BIV. In such possibilities, she would still believe that she is not a BIV, so her belief is insensitive, as desired.

At this point, one may think that, given that *Method Sensitivity* and *Method Safety* only state necessary conditions for knowledge, proponents of these accounts could appeal to some other non-modal conditions to explain Suzy's ignorance. However, I think this option is unattractive. This is because the deficiency involved in Suzy's belief is precisely a deficiency of a modal nature—a deficiency that stems from the close error-possibilities of forming false beliefs. Indeed, we may simply add some details to Suzy's case so that her belief that she is not a BIV, just like our belief that we are not BIVs, satisfies any non-modal conditions. It then follows that it is a desideratum of a modal account like sensitivity or safety to be able to explain Suzy's ignorance. In other words, the kind of deficiency involved in Suzy's belief *should* be explained by such modal accounts, as opposed to other non-modal conditions.

4 An Objection

Now, let us consider a potential objection. In particular, one may bring up the familiar distinction between knowledge-friendly 'evidential luck' and knowledge-precluding 'veritic luck' (Cf. Engel 1992; Pritchard 2005), and argue that Suzy's case only involves the former. Roughly speaking, a belief is veritically lucky when it is *true* as a matter of luck. By contrast, a belief is evidentially lucky when a subject is lucky to be in her evidential situation, but given the situation she is in, it is not a matter of luck that the belief is true. For instance, one is lucky to see a meteor shooting across the sky when looking out of the window at the right moment. In this sense, her belief that a meteor just shot across the sky is evidentially lucky. But such luck does not seem to prevent the person from knowing. Likewise, one may think that in Suzy's scenario, she is lucky to be in the evidential situation she is in—a situation under which she enjoys veridical experiences. But once she is in such evidential situation, it is not a matter of luck that she forms a true belief that she is not a BIV. Considered this way, it may be argued that the kind

[8] Incidentally, compare the possibility of Suzy becoming a BIV and the possibility of dreaming. One commonality between these possibilities is that they are both close ones. The difference, however, is that if Suzy were a BIV she could still form beliefs (via the hallucinatory experiences instilled to her by the scientists). By contrast, as Sosa (2007, chapter 1) argues, while dreaming we are not really believing, but merely imagining. If so, the possibility of dreaming does not have implications about how belief-formation methods should be individuated, since dreaming possibilities would be disregarded as irrelevant ones for the purpose of evaluating sensitivity and safety of a belief no matter how methods are being individuated.

of luck involved is of the evidential variety, so that Suzy's belief *does* amount to knowledge. Perhaps the intuition that she does not know only emerges from a confusion between evidential luck and veritic luck.

I have three responses to this objection. First, notice that a belief being evidentially lucky does not imply that it is not veritically lucky. After all, these forms of luck are not mutually exclusive: It is possible that both forms of luck are involved in a belief. So, even if Suzy's case involves knowledge-friendly evidential luck, it does not entail that vertic luck is absent.

Second, Suzy's case is not strictly analogous to the meteor scenario: in the former there are many close error-possibilities in which Suzy has the same experiences but her belief is false. In comparison, there are no such close error-possibilities in the meteor case. In this regard, Suzy's situation is more analogous to the cases that involve environmental luck, which is a form of veritic luck that results from one's being in an unfriendly epistemic environment. To give an example of this form of luck, suppose that you happen to see the only zebra in a zoo where, unbeknownst to you, numerous cleverly disguised mules that look perfectly like a zebra stand in the vicinity (Cf. Dretske 1970; Goldman 1976). As many concur, in such a case you do not know that you are facing a zebra. This is analogous to Suzy's situation in which she believes that she is not a BIV based on, say, her experiences, but there are many close possibilities in which she would be misled by the scientists and falsely believe that she is not a BIV. Both situations involve unfriendly epistemic environment, so that in both cases there are close error-possibilities in which the subject is misled by undistinguishable experiences and forms false beliefs. Thus considered, the kind of luck involved in Suzy's belief *is* akin to the knowledge-undermining environmental kind.

Third, and relatedly, consider Pritchard's much-discussed definition of veritic luck:

> Veritic Luck (VL): S's true belief is lucky *iff* there is a wide class of near-by possible worlds in which S continues to believe the target proposition, and the relevant initial conditions for the formation of that belief are the same as in the actual world, and yet the belief is false (Pritchard 2007, p. 280).

Whether Suzy's belief features (environmental-) veritic luck according to VL substantially depends on what count as 'the relevant initial conditions' of Suzy's belief-formation. If one insists that the conditions in her actual situation and the conditions in the counterfactual situations where she is a BIV are different, then VL would not rule Suzy's belief as being veritically lucky. Notice, however, that in order to predict that in the zebra case the subject's belief *is* vertically lucky, one has to individuate initial conditions in a way that such conditions essentially incorporate the experiences associated with one's actual belief. If not, in the zebra case the close error-possibilities in which the subject sees cleverly disguised mules and forms false beliefs would be counted as irrelevant, so that one cannot predict that the actual belief is vertically lucky. But if relevant initial conditions do incorporate one's experiences, presumably in Suzy's case there are many close error-possibilities in which the initial conditions of her belief-formation are the same (i.e., close BIV worlds), so that her actual belief that she is not a BIV *is* vertically lucky.

Finally, let me address some lingering worries one may have with regards to my third reply. To start with, one may argue that there are still some relevant differences between Suzy's case and the zebra case. For instance, one may argue that in the former, when Suzy is turned into a BIV she suffers from major bodily changes. In the zebra case, however, when the subject forms false belief in close error-possibilities she does not suffer from such changes. Therefore, the relevant initial conditions in the actual situation and the conditions in the close error-possibilities are still different in these two cases. However, I do not think the difference here is epistemically relevant. For we could simply stipulate that in the zebra case, when the subject perceives those cleverly disguised mules, she would also suffer from major bodily changes (for whatever reasons), but this hardly changes the fact that this case involves knowledge-undermining luck.

Besides, one may argue that in the zebra case there are indistinguishable but misleading objects located in the vicinity of the subject's actual physical environment, whereas in Suzy's case, there are no such objects around. In this sense, one may think that in Suzy's case (but not in the zebra case), the epistemic environment is only *potentially* unfriendly, not actually unfriendly. And perhaps this difference matters for whether or not veritic luck is involved in Suzy's case. Here, again, I think such a difference does not carry much weight. For one thing, if one thinks that whether or not there are indistinguishable objects located in one's environment really matters, this is a point that requires further argument. After all, in general, what modal epistemologists really care about is whether or not the truth of a belief is robust throughout certain modal space, not the specific mechanism that results in such robustness (or lack thereof). For another, the distinction between 'potentially unfriendly environment' and 'actually unfriendly environment' is not sharp. In a sense, we could say that Suzy's environment is actually unfriendly—there are criminals who turn people into BIVs in the vicinity. Furthermore, in the zebra case, given that the subject does not actually perceive cleverly disguised mules, there is a sense in which the environment is only potentially unfriendly for the subject—that is, she

potentially could have seen those mules, but does not actually perceive them anyway.[9]

5 Implications

In closing, let us glean some general lessons from the above discussions. First, any account that insists that we know we are not BIVs may face a similar challenge as the modal accounts that we criticized above, provided that such an account's explanation of why we know hinges on a feature that is shared by Suzy's scenario. For instance, consider Kelp's (2021) recent reply to the skeptic. As mentioned earlier, Kelp argues that we may believe, and even know, that we are not BIVs through an epistemic ability of recognizing that BIV possibilities could not easily obtain. Plausibly, such an ability can be shared by someone like Suzy as well[10] (though in her case this ability can only generate accidentally/luckily true belief that she is not a BIV, as opposed to knowledge). It then follows that on Kelp's account, Suzy's true belief also constitutes knowledge, implausibly. The lesson is that a theorist who wants to insist that we have anti-skeptical knowledge must offer an explanation of such knowledge that cannot be carried over to Suzy's situation, on pain of attributing such knowledge even when skeptical possibilities are rather close.

Second, for sensitivity theorists, they have to reconsider whether or not they should adopt a neo-Moorean reply to the sceptic, like the one Black (2002) endorses. Black's approach to methods generates the result that I know that I am not a BIV. This means that contra Nozick's original account, Black's approach helps a sensitivity theorist to better accommodate the epistemic closure principle. This may strike some as an advantage of Black's account over Nozick's. But now, it becomes clear that Nozick's account has some advantage that Black's approach lacks: The latter incurs the cost that one can know that one is not a BIV even when this is a very close possibility, whereas Nozick's account does not imply this much.[11]

Finally, proponents of modal accounts of knowledge have to rethink about their theories in the context of internalism/externalism distinction. Typically, both sensitivity and safety are broadly construed as externalist accounts of knowledge, in the sense that whether a belief is safe or sensitive depends on the belief's external relation to the world—a relation that the subject does not need to have any access to (Pritchard 2008). However, if my above argument is sound such that modal theorists should embrace *SESM*, there would at least be some internalist element in sensitivity and safety. This is because *SESM* states that methods should be construed

[9] An anonymous reviewer suggests that, instead of the earlier meteor case, a case of evidential luck that is more analogous to Suzy's case is the following. A group of students is broken into 8 groups via lottery. The students are supposed to learn about topic X. Lucy was assigned to group 1, which luckily for her, is the only group that has a teacher worth their salt. Students in the other groups will acquire a bunch of false beliefs and no knowledge about X, but students in Lucy's group will come to know everything there is to know about X. It seems clear this case only involves evidential luck. But then, shouldn't we say the same thing about Suzy's case?

In reply, although I do agree that this case seems more analogous to Suzy's case, it does not show that Suzy's belief is immune from veritic (environmental) luck. Again, there are close possible worlds in which Suzy has the same but misleading experiences, but this feature does not obtain in the above case: if Lucy were assigned to a different group, she would have different experiences by encountering different teachers who spread falsehoods about topic X. This implies that, on Pritchard's definition of veritic luck, Lucy's belief should not be counted as a case of veritic luck, as the initial conditions in the actual and the counterfactual situations are different. By contrast, Suzy's belief could plausibly be counted as a case of veritic luck, as argued earlier. Apart from that, I think there is a sense in which Suzy's local environment is unfriendly: there are criminals who are ready to turn people into BIVs anytime in the vicinity of Suzy's location. The threat is always there, so to speak. By contrast, in Lucy's case, as long as she is not assigned to the other groups, she stays in a friendly local environment. I think this partly explains why the intuition of Suzy's ignorance is so strong that it cannot be explained away by the similarities between Suzy's case and Lucy's case.

[10] An anonymous reviewer helpfully points out that this depends on whether 'recognizing' is factive. If it is factive, then someone like Suzy does not have the ability to recognize that BIV possibilities are remote. I have three replies. First, although Kelp does not discuss the issue of whether or not recognizing is factive, he points out that there are two important properties of beliefs produced by recognitional abilities (Kelp 2021, p. 382). First, they are spontaneous—that is, they are not based on further evidence. Second, people are often unable to articulate how their recognitional abilities work in detail. Now, if our beliefs that we are not BIVs (which, as Kelp argues, are produced by recognitional abilities) have these two properties, plausibly Suzy's belief that she is not a BIV could share these properties as well, such that it is a good candidate of a recognitional belief pro-

Footnote 10 (continued)

duced by her recognitional abilities. After all, we may stipulate that Suzy is an internal duplicate of an agent in our world, so that just as the agent in our world is unable to articulate how the recognitional abilities work in detail and that her recognitional belief is not based on any further evidence, Suzy's belief has these properties as well. Second, I take it as relatively uncontroversial that, at least according to the ordinary usage of the term, 'recognizing' is non-factive. For instance, it makes sense to say that "I recognized that guy as Tom, but it turns out that he is Tom's twin brother." Finally, if one has reservations about whether recognizing is non-factive or not, the present point can be simply taken as a conditional claim that *if* recognizing is non-factive, *then* an account like Kelp's cannot explain Suzy's ignorance.

[11] Apart from the present objection against Black's sensitivity account, Melchior (2015, 2019) argues that the account cannot properly accommodate higher-level knowledge. That is, according to the account, one can know that p without knowing that one does not falsely believe that p. In addition, Melchior argues that Black's account permits that we know that we are not BIVs but it predicts that other beliefs formed via bootstrapping are insensitive (such as my belief that my gas gauge is reliable).

in an internalist fashion so that when two method tokens are experientially the same, they must be of the same type of method, regardless of any external differences involved, just as Nozick (1981) argues. Interestingly, this echoes with some recent process reliabilists' move. Such reliabilists as Goldman (2011) and Comesaña (2010) have proposed hybrid accounts which combine reliabilism with evidentialist elements (Cf. Alston 1988; Tang 2016). Now it seems that sensitivity and safety theorists should move towards a similar direction.

Acknowledgements For helpful comments and discussions, I'd like to thank Ge Fang, Xingming Hu, Maggie Huang, Xiaoyu Ke, Xuya Ma, Masahiro Yamada, Chong Yuan, Yiling Zhou, audiences of 10th Asian Epistemology Network meeting, and two anonymous reviewers for this journal.

Funding This research was funded by the "Young Scholar Project in Humanities and Social Sciences" of China's Ministry of Education (Grant Number: 21YJC720022) and National Social Science Fund of China (Grant Number: 22CZX016).

References

Alfano M (2009) Sensitivity theory and the individuation of belief-formation methods. Erkenntnis 70(2):271–281
Alston W (1988) An internalist externalism. Synthese 74(3):265–283
Becker K (2007) Epistemology modalized. Routledge, London
Becker K (2008) Epistemic luck and the generality problem. Philos Stud 139(3):353–366
Becker K (2012) Methods and how to individuate them. In: Becker K, Black T (eds) The sensitivity principle in epistemology. Cambridge University Press, New York, pp 81–97
Beddor B, Pavese C (2020) Modal virtue epistemology. Philos Phenomenol Res 101:61–79
Black T (2002) A Moorean response to brain-in-a-vat scepticism. Australas J Philos 80(2):148–163
Black T (2008) Defending a sensitive neo-Moorean invariantism. In: Hendricks V, Pritchard D (eds) New waves in epistemology. Palgrave-Macmillan, London, pp 8–27
Black T, Murphy P (2007) In defense of sensitivity. Synthese 154(1):53–71
Bogardus T, Marxen C (2014) Yes, safety is in danger. Philosophia 42(2):321–334
Broncano-Berrocal F (2014) Is safety in danger? Philosophia 42(1):1–19
Broncano-Berrocal F (2018) Knowledge and tracking revisited. Analysis 78(3):396–405
Comesaña J (2010) Evidentialist Reliabilism. Noûs 44(4):571–600
Cross T (2010) Skeptical success. Oxford Studies in Epistemology 3:35–62
DeRose K (2010) Insensitivity is back, baby! Philos Perspect 24(1):161–187
Dretske F (1970) Epistemic operators. J Philos 67(24):1007–1023
Engel M (1992) Is epistemic luck compatible with knowledge? South J Philos 30(2):59–75
Goldman A (1976) Discrimination and perceptual knowledge. J Philos 73(11):771–791
Goldman AI (2011) Toward a synthesis of reliabilism and and evidentialism? Or: evidentialism's troubles, reliabilism's rescue package. In: Dougherty T (ed) Evidentialism and its discontents. Oxford University Press, pp 254–280
Greco J (2016) Knowledge, virtue and safety. In: Ángel M (ed) Performance epistemology. Oxford University Press, New York, pp 51–61
Grundmann T (2020) Saving safety from counterexamples. Synthese 197(12):5161–5185
Hirvelä J (2019) Global safety: how to deal with necessary truths. Synthese 196(3):1167–1186
Kelp C (2021) Theory of inquiry. Philos Phenomenol Res 103(2):359–384
Melchior G (2015) The heterogeneity problem for sensitivity accounts. Episteme 12(4):479–496
Melchior G (2019) Knowing and checking: an epistemological investigation. Routledge, London
Nozick R (1981) Philosophical explanations. Harvard University Press, Cambridge
Peet A, Pitcovski E (2018) Normal knowledge: toward an explanation-based theory of knowledge. J Philos 115(3):141–157
Pritchard D (2005) Epistemic luck. Oxford University Press, Oxford
Pritchard D (2007) Anti-luck epistemology. Synthese 158(3):277–297
Pritchard D (2008) Sensitivity, safety, and anti-luck epistemology. In: Greco J (ed) The Oxford handbook of skepticism. Oxford University Press, Oxford, pp 437–455
Pritchard D (2012) Anti-luck virtue epistemology. J Philos 109(3):247–279
Roush S (2007) Tracking truth: knowledge, evidence, and science. Oxford University Press, Oxford
Smith M (2016) Between Probability and Certainty: What Justifies Belief. Oxford University Press, Oxford
Sosa E (1999) How to defeat opposition to Moore. Philos Perspect 13:141–154
Sosa E (2007) A virtue epistemology: apt belief and reflective knowledge. Oxford University Press, Oxford
Tang WH (2016) Reliability theories of justified credence. Mind 125(497):63–94
Williamson T (2000) Knowledge and its limits. Oxford University Press, Oxford
Williamson T (2009) Reply to Goldman. In: Pritchard D, Greenough P (eds) Williamson on knowledge. Oxford University Press, Oxford, pp 305–312
Zalabardo J (2017) Safety, sensitivity and differential support. Synthese 197(12):5379–5388