Sensitivity, Safety, and Epistemic Closure

Bin Zhao

Abstract It has been argued that an advantage of the safety account over the sensitivity account is that the safety account preserves epistemic closure; while the sensitivity account implies epistemic closure failure. However, the argument fails to take the method-relativity of the modal conditions on knowledge, viz, sensitivity and safety, into account. In this paper, I argue that the sensitivity account and the safety account are on a par with respect to epistemic closure once the method-relativity of the modal conditions is taken into account. Therefore, epistemic closure is no longer an arbiter in the debate.

1. Introduction

Modal epistemologies maintain that there must be a modal connection between one's belief and the truth for the belief to count as knowledge: one not only should hold a true belief in the actual world, but also should avoid holding a false belief across a specified range of nonactual possible worlds. Among the main modal epistemologies in the literature are the

sensitivity account of knowledge[1] and the safety account of knowledge[2].

Though both accounts handle a wide range of cases involving knowledge-precluding epistemic luck as well as cases of knowledge,[3] it has been argued that an advantage of the safety account over the sensitivity account is that the safety account preserves epistemic closure; while the sensitivity account implies epistemic closure failure. In this paper, I demonstrate that such an argument fails to take into account that the modal conditions are usually relativized to belief-formation methods. Once we take the method-relativity of the modal conditions into account, it should be clear to us that the sensitivity account and the safety account are on a par with respect to epistemic closure, viz., either both accounts preserve epistemic closure or both accounts imply epistemic closure failure. Therefore, epistemic closure cannot be used to adjudicate between the sensitivity and the safety accounts of knowledge.[4]

---

[1] The sensitivity theorists include Adams et al. (2012), Adams & Clarke (2005), Becker (2006, 2007, 2008, 2009, 2012a, 2012b, 2013, 2016, 2018), Bjerring & Gundersen (2020), Black (2002, 2008, 2011, 2019), Black & Murphy (2007), Bolos & Collin (2018), Broncano-Berrocal (2018), Collin (2018), Collin & Bolos (2020), Cross (2010), DeRose (1995, 2010), Goldberg (2012), Gundersen (2003, 2010, 2012), Ichikawa (2011), Murphy & Black (2012), Nozick (1981), Ramachandran (2015), Roush (2005, 2012), Topey (2021), Wallbridge (2017, 2018a, 2018b), and Zalabardo (2012, 2017).

[2] The safety theorists include Ball (2016), Beddor and Pavese (2020), Carter (2010), Dutant (2010, 2016), Greco (2012), Grundmann (2020), Hirvelä (2017, 2019), Kelp (2013), Lasonen-Aarnio (2010), Luper (2003, 2006a, 2006b), Luper-Foy (1984), Manley (2007), Neil (2021), Peet and Pitcovski (2018), Pritchard (2002, 2005, 2007, 2008a, 2008b, 2008c, 2009, 2012a, 2012b, 2013, 2015, 2016, 2017, 2018), Sainsbury (1997), Wedgwood (2020), and Williamson (2000, 2009). In addition, Sosa (1999a, 1999b, 2003, 2015) was also a main proponent of the safety condition though he has recently developed a virtue-theoretic account of knowledge that does not draw upon the condition heavily, see Sosa (2007, 2009).

[3] Nonetheless, it has been argued that the sensitivity account, unlike the safety account, fails to account for inductive knowledge, See Pritchard (2008b) and Sosa (1999b). For the argument that the sensitivity account is able to account for inductive knowledge, see Wallbridge (2018). For the argument that the safety account also fails to account for inductive knowledge, see Zhao & Baumann (2021).

[4] Bernecker (2012) also argues that epistemic closure cannot be used to adjudicate between the safety and the sensitivity accounts of knowledge. As he argues, sensitivity, as well as safety, fails to account for knowledge of necessary truths which is required in competent deductions. Because whether the safety/sensitivity account can account for knowledge of necessary truths bears little relevance to the topic of the paper, I shall not delve into his argument here. For related discussions, see Ball (2016), Bernecker (2020), Blome-Tillmann (2017), Collin (2018), Garrett (1983), Grundmann (2020), Hirvelä (2017, 2019), Manley (2007), Melchior (2017a, 2021), Miščević (2007), Paterson (2020), Pritchard (2009, 2012a, 2013, 2016), Roland & Cogburn (2011), Zhao (2021a, 2021b, forthcoming), and Zhao (2021).

2. Sensitivity and Safety

Consider a scenario where one looks at a reliable clock in normal lighting conditions and thus forms a true belief that *it is now 12:00*. Since the clock is reliable and the lighting conditions are normal, the belief counts as knowledge. Consider another scenario where one looks at a clock that stops 12 hours earlier in normal lighting conditions and thus forms a true belief that *it is now 12:00*. The belief is true as a matter of luck and thus does not count as knowledge.

The lesson is that luckily true beliefs do not count as knowledge. To eliminate luckily true beliefs from the realm of knowledge, an anti-luck condition on knowledge is needed. A natural idea is that if one's belief is true as a matter of luck, then one would still believe the same thing even if it turns out to be false. As Murphy and Black write,

> "How, then, do we keep luckily acquired beliefs from counting as knowledge? We must demand more of S than that she respond appropriately to her environment by suitably forming the true belief that P. One idea is to demand that S respond appropriately in environments where it is *not* the case that P" (Murphy & Black, 2012, p. 30).

This idea motivates the sensitivity account of knowledge.[5] According to this account, *S* knows that *p* only if *S*'s belief in *p* is sensitive, that is, *S* would not believe that *p* if *p* were false, or formally $\sim p \Rightarrow \sim Bp$ ("$\Rightarrow$" denotes the subjunctive conditional connective) (Nozick, 1981, p. 177).

The sensitivity condition is rendered as follows under the Lewis-Stalnaker analysis of

---

[5] For discussions of the sensitivity condition as the anti-luck condition on knowledge, see Adams & Clarke (2005), Becker (2007, 2008, 2012b), Black (2011, 2019), Bolos & Collin (2018), Collin (2018), Murphy & Black (2012), Neil (2021), Roush (2005), and Topey (2021).

subjunctive conditionals (Lewis, 1973a, 1973b; Stalnaker, 1968),

> **SENSITIVITY**: S's belief in p is sensitive just in case, in the closest possible world where p is false, S does not believe that p.

This makes us consider whether S holds a false belief in p in the closest possible world where p is false. If S believes that p in the possible world, then S's belief in p is insensitive. Thus, the belief is luckily true and S does not know that p. If S does not believe that p in the possible world, then S's belief in p is sensitive. Thus, the belief is non-luckily true and S knows that p unless it exhibits some non-modal shortcomings that would deprive it of the status of knowledge.

Another idea is that if one's belief is true as a matter of luck, then the belief could easily have been false. This idea motivates the safety account of knowledge.[6] According to this account, S knows that p only if S's belief in p is safe, that is, only if S could not easily have falsely believed p, or formally $Bp \Rightarrow p$.

The safety condition is usually rendered as follows,

> **SAFETY**: S's belief in p is safe just in case, in all nearby possible worlds where S forms a belief that p, p is true.

This makes us consider whether p is true in nearby possible worlds where S believes that p. If p is false in some of these possible worlds, then S's belief in p is not safe. Thus, the belief is luckily true and S does not know that p. If p is true in all these possible worlds, then S's belief

---

in *p* is safe. Thus, the belief is non-luckily true and *S* knows that *p* unless it exhibits some non-modal shortcomings that would deprive it of the status of knowledge.

Both SENSITIVITY and SAFETY produce a welcome result in the above scenarios. For SENSITIVITY, one knows the correct time in the first scenario because one's belief is sensitive. In the closest possible world where it is not 12:00, i.e., the possible world where one looks at the clock one minute earlier or later, the reliable clock would not indicate 12:00, thus one does not believe that *it is now 12:00* by reading the reliable clock. One does not know the correct time in the second scenario because one's belief is insensitive which accounts for why it is true as a matter of luck. In the closest possible worlds where it is not 12:00, i.e., the possible world where one looks at the clock one minute earlier or later, the stopped clock still indicates 12:00, thus one believes that *it is now 12:00* by reading the stopped clock. For SAFETY, one knows the correct time in the first scenario because one's belief is safe. In all nearby possible worlds where one believes that *it is now 12:00* on the clock, it is 12:00. One does not know the correct time in the second scenario because one's belief is unsafe. There are some nearby possible worlds where he looks at the clock one minute earlier and believes that *it is now 12:00* on the clock, while the time is 11:59. In a word, both accounts handle cases involving knowledge-precluding epistemic luck as well as cases of knowledge nicely.

However, both SENSITIVITY and SAFETY run into counterexamples quickly. For instance,

DACHSHUND: "I look at a nearby dachshund and truly believe that what I see is a dog. Had I not been seeing a dachshund, I would have been seeing a wolf, and would have

falsely believed myself to be seeing a dog" (Goldman, 2009, p. 80).[7]

Intuitively, I know *what I see is a dog* despite my falsely believing so in the closest possible world where what I see is not a dog, i.e., what sees in that possible world is a wolf. In addition, there is a nearby possible world where I falsely believe that *what I see is a dog.* Thus, we have a counterexample for both SENSITIVITY and SAFETY.

Once it is noted that the wolf does not look like the dog, we may want to relativize the sensitivity condition to the belief-formation method to get rid of the counterexample. If the condition is relativized in that way, then a possible world should be taken into account only if it is a possible world where S uses the same method as that in the actual world; while possible worlds where S uses a different method are irrelevant. That is to say, which possible worlds are relevant depends on the individuation of belief-formation methods. In DACHSHUND, the belief-formation methods are individuated by reference to perceptual appearances. Since the wolf appears to be different from the dog, the subject's belief in the counterfactual case is formed on a different method from that in the actual case (Goldman, 2009, p. 81). Therefore, the subject's belief is sensitive if the sensitivity condition is relativized to the belief-formation methods.[8] The sensitivity condition is thus rendered as follows,

> **SENSITIVITY^M**: S's belief in p which is formed on method M is sensitive just in case,
>
> in the closest possible world where p is false and S uses M to form a belief whether or not

---

[7] Similar counterexamples include Alfano's (2009) REDWOOD, Goldman's (1976) JUDY & TRUDY, and Nozick's (1981) GRADNMOTHER.

[8] In addition to discharging putative counterexamples such as DACHSHUND, there are other motivations for relativizing the sensitivity condition to belief-formation methods. For instance, there is criticism of the sensitivity account that it cannot account for higher-order knowledge, see DeRose (1995), Huemer (2001), Kripke (2011), Melchior (2015, 2017b), Sosa (1996, 1999b, 2002), Vogel (1987, 2000, 2007, 2012), Williamson (2000), and Zalabardo (2012). Nonetheless, Bjerring & Gundersen (2020) and Wallbridge (2017, 2018a) argue that the criticism fails to take the method-relativity of the sensitivity condition into account.

p, *S* does not believe that *p*.

Similarly, the safety is usually relativized to belief-formation methods to avoid counterexamples such as DACHSHUND. The safety condition is thus rendered as follows,

SAFETY^M: *S*'s belief that *p*, formed on method *M*, is safe, if and only if, in all nearby possible worlds where *S* forms a belief that *p* on method *M*, *p* is true.

If the safety condition is thus relativized to belief-formation methods, then counterexamples such as DACHSHUND would be discharged similarly.

3. The Consideration of Epistemic Closure

Though both accounts handle a wide range of cases involving knowledge-precluding epistemic luck as well as cases of knowledge, it has been argued that the sensitivity account implies epistemic closure failure (Kripke, 2011; Luper, 2012; Pritchard 2002, 2005, 2008b; Sosa, 1999b, 2004; Vogel, 1987; Williamson, 2000); while the safety account preserves epistemic closure (Luper, 2012; Pritchard 2002, 2005, 2008b; Sosa, 1999b; 2004). Thus, the consideration of epistemic closure is in favor of the safety account.[9]

Epistemic closure is the idea that knowledge is closed under competent deduction in the sense that we can extend our knowledge by competently deducing consequences from what is already known. For instance, if I know that *I am looking at a red barn*, then I can know that *I am looking at a barn* by deduction. To put it formally,

Epistemic Closure: If *S* knows *p* and competently deduces *q* from *p*, thereby coming to

---

[9] For dissenters who think that the safety account also implies epistemic closure failure, see Alspector-Kelly (2011), Kvanvig (2008b), and Murphy (2005, 2006).

believe *q*, while retaining *S*'s knowledge that *p*, *S* comes to know that *q* (Hawthorne, 2005, p. 29).[10]

This principle seems to be a plain truth to the effect that many epistemologists take it for granted without further arguments. For instance, Steven Luper writes that "virtually everyone agrees that the closure principle, suitably formulated, is highly plausible" (Luper, 2012, p. 222). After all, what can be plainer than knowing that *I am looking at a barn* by deducing it from my knowledge that *I am looking at a red barn*?

Nonetheless, it has been argued that *S* can have an insensitive belief in *q*, which is competently deduced from *p* which *S* knows to be true. For instance, I know that *I have hands* because my belief in it is sensitive, viz., in the closest possible world where I do not have hands, e.g., a possible world where I lost my hands in a severe car accident, I do not believe that *I have hands*. However, I do not know that *I'm not a handless brain-in-a-vat* because my belief in it is insensitive, viz., in the closest possible world where I'm a handless brain-in-a-vat, I still believe that *I'm not a handless brain-in-a-vat*, though *I'm not a handless brain-in-a-vat* is competently deduced from *I have hands* which I know to be true. Therefore, the sensitivity account implies epistemic closure failure (Pritchard 2002, 2005, 2008b; Sosa 1999b; Vogel, 1987; Williamson, 2000).

In turn, let us see how the safety account performs when it comes to epistemic closure. According to the safety account, I know that *I have hands* because my belief in it is safe, viz., in nearby possible worlds where I believe that *I have hands*, it is true that I have hands

---

[10] This formulation needs some further qualifications. For example, *S* learns of no undefeated defeater for *q*. For the sake of simplicity, I shall avoid further complications here. For related discussions, see David and Warfield (2008), Kvanvig (2006), Luper (2020), Warfield (2004).

because possible worlds where I falsely believe so, e.g., the possible worlds where I'm a handless brain-in-a-vat, are very remote. I also know that *I'm not a handless brain-in-a-vat* because my belief in it is safe, viz., in nearby possible worlds where I believe that *I'm not a handless brain-in-a-vat*, it is true that I'm not a handless brain-in-a-vat given that the possible worlds where I'm a handless brain-in-a-vat are very remote. Thus, it has been argued that the safety account is better placed than the sensitivity account to preserve epistemic closure which, in turn, is one of the main reasons in favor of the safety account.

4. The Method-Relativity of Modal Conditions and Epistemic Closure

However, the aforementioned argument does not pay enough attention to the method-relativity of modal conditions. In a case where S competently deduces q, e.g., *I'm not a handless brain-in-a-vat*, from p, e.g., *I have hands*, which S knows to be true. S not merely believes that q, but believes it on some belief-formation method. When evaluating whether the belief that q is sensitive or not, the argument examines the closest possible world where the target proposition is false, e.g., the closest possible world where I'm a handless brain-in-a-vat. Nonetheless, what we should examine is the closest possible world where the target proposition is false and the subject uses the same method as that in the actual world. It is unclear whether S's belief that q is sensitive or not once the method-relativity of the modal conditions is taken into account.

There are different ways to individuate the belief-formation method: it could be individuated as inference from the belief that p (Alspector-Kelly, 2011; Nozick, 1981), or individuated as inference from the justified belief that p, or individuated as inference from the

true belief that *p*; or individuated as inference from knowledge that *p* (Hirvelä, 2020; Schulz, 2020; Williams, 1991), etc.[11] Depending on whether the truth of *p* is ensured by the individuation, the individuations could be classified into two sorts: the individuations that do not ensure the truth of *p*, e.g., inference from the belief that *p* or inference from the justified belief that *p*. If the method is thus individuated, then one is able to use the method no matter whether *p* is true or not; and the individuations that ensure the truth of *p*, e.g., inference from the true belief that *p* or inference from knowledge that *p*. If the method is thus individuated, then one is able to use the method only if *p* is true.

Thus, which possible world is relevant when evaluating whether the target belief is sensitive or not depends on how we individuate the belief-formation method of the target belief. How do we determine which individuation is the correct one? We seem to encounter the generality problem here. Instead of committing myself to a specific individuation of the belief-formation methods, I shall argue that the sensitivity account and the safety are on a par with respect to epistemic closure no matter whether the belief-formation method is individuated such that the truth of *p* is ensured or not.

4.1 Individuations that Do Not Ensure the Truth of *p*

In a case where *S* competently deduces *q* from *p* which *S* knows to be true, if the belief-formation method is individuated such that the truth of *p* is not ensured, e.g., inference from the belief that *p* or inference from the justified belief that *p*, then it is the case that *S*'s belief that *q* could be both insensitive and unsafe.

---

[11] For discussions of the individuation of belief-formation methods, see Alfano (2009), Becker (2008, 2009, 2012b), Bernecker (2018), Black & Murphy (2007), Bogardus & Marxen (2014), Broncano-Berrocal (2014), Hirvelä (2019), Zhao (2021a), and Zhao (2020).

For the belief to be insensitive, *S* needs to believe that *q* in the closest possible world where *q* is false and *S* uses the same method as that in the actual world, e.g., inference from the belief that *p* or inference from the justified belief that *p*, to form a belief whether or not *q*. Since *p* entails *q*, *p* is also false in that possible world. In addition, *S* should believe that *p* in that possible world. Otherwise, *S* can not infer *q* from *p*. Thus, *S* believes that *p* in a possible world where *p* is false. Since *S* knows that *p*, *S*'s belief that *p* in that possible world should be formed on a different method from that in the actual world or that possible world is not the closest possible world where *p* is false and *S* uses the same method as that in the actual world. As long as this desideratum is satisfied, we would have an epistemic closure failure case for the sensitivity theorists.

For the belief that *q* to be unsafe, *q* needs to be false in some nearby possible world where *S* believes that *q* on the same method as that in the actual world, e.g., inference from the belief that *p* or inference from the justified belief that *p*. Since *p* entails *q*, *p* is also false in that possible world. In addition, *S* should believe that *p* in that possible world. Otherwise, *S* can not infer *q* from *p*. Thus, *p* is false in some nearby possible world where *S* believes that *p*. Since *S* knows that *p*, *S*'s belief that *p* in that possible world should be formed on a different method from that in the actual world. As long as this desideratum is satisfied, we will have an epistemic closure failure case for the safety theorists.

A revised version of Goldman's DACHSHUND satisfies the aforementioned desideratum,

DACHSHUND II: I look at a nearby dachshund and truly believe that *what I see is a dog*. Had I not been seeing a dachshund, I would have been seeing a wolf, and would

have falsely believed myself to be seeing a dog. I then infer that what I see is not a wolf.[12]

Intuitively, I know *what I see is a dog*. What about my inferential belief that *what I see is a wolf*? In the closest possible world where *what I see is a wolf* and I use the same method as that in the actual world, e.g., inference from the belief that *what I see is a dog* or inference from the justified belief that *what I see is a dog*, to form a belief whether what I see is a wolf or not, I believe *what I see is not a wolf*. Thus, my belief that *what I see is not a wolf* is insensitive. In addition, in some nearby possible world where I believe *what I see is not a wolf* on the same method as that in the actual world, e.g., inference from the belief that *what I see is a dog* or inference from the justified belief that *what I see is a dog*, what I see is a wolf. Thus, my belief that *what I see is not a wolf* is unsafe. Therefore, my belief that *what I see is not a wolf* is both insensitive and unsafe and thus falls short of knowledge though it is competently deduced from what I know to be true.

Here is a recipe to construct an example that shows that the sensitivity account and the safety account imply epistemic closure failure: First, construct a case where *S* competently deduces *q* from *p* which *S* knows to be true on some belief-formation method. Second, make sure that there is a possible world where *S* falsely believes that *p* on a different method. Third, make sure that *S* falsely believes that *q* by competent deduction from *S*'s belief that *p* in that possible world. Forth, make sure that the possible world is a nearby one and it is also the closest possible world where *q* is false. Once these conditions are satisfied, we will have an example showing that both accounts imply epistemic closure failure.

---

[12] Similar examples include Alspector-Kelly's (2011) LAMBORGHINI and Kripke's (2011) RED BARN.

In a word, both accounts imply epistemic closure failure if the belief-formation method is individuated such that the truth of p is not ensured.

4.2 Individuations that Ensures the Truth of p

In a case where S competently deduces q from p which S knows to be true, if the belief-formation method is individuated such that the truth of p is ensured, e.g., inference from the true belief that p or inference from knowledge that p, then it is the case that S's belief that q is both insensitive and unsafe.

If the belief-formation method is thus individuated, then p is true in all possible worlds where S forms a belief that q on the same method as that in the actual world. Since p logically entails q, q is also true in these possible worlds. Thus, S's belief that q is both sensitive and safe because there is no possible world where S falsely believes that q on the same method as that in the actual world.

In a word, both accounts preserve epistemic closure if the belief-formation method is individuated such that the truth of p is ensured.

5. Conclusion

Whether the sensitivity account and the safety account preserve epistemic closure or not depends on how the belief-formation method of inferential beliefs is individuated. In a case where S competently deduces q from p which he knows to be true. If the method is individuated such that the truth of p is ensured, e.g., inference from the true belief that p or inference from knowledge that p, then S's belief that q is both sensitive and safe; if the

method is individuated such that the truth of *p* is not ensured, e.g., inference from the belief that *p* or inference from the justified belief that *p*, then *S*'s belief that *q* could be both insensitive and unsafe. The conclusion is that either both accounts preserve epistemic closure or both accounts imply epistemic closure failure. Therefore, the consideration of epistemic closure does not help us to advance the debate in modal epistemology.[13]

REFERENCES

Adams, Fred, Barker, John A., & Figurelli, Julia. (2012). "Towards Closure on Closure." *Synthese*, 188 (2): 179–196.

Adams, Fred, & Clarke, Murray. (2005). "Resurrecting the Tracking Theories." *Australasian Journal of Philosophy*, 83 (2): 207–221.

Alfano, Mark. (2009). "Sensitivity Theory and the Individuation of Belief-Formation Methods." *Erkenntnis*, 70 (2): 271–281.

Alspector-Kelly, Marc. (2011). "Why Safety Doesn't Save Closure." *Synthese*, 183 (2): 127–142.

Ball, Brian. (2016). "Knowledge, Safety, and Questions." *Filosofia Unisinos*, 17 (1), 58–62.

Becker, Kelly. (2006). "Is Counterfactual Reliabilism Compatible with Higher-Level Knowledge?" *Dialectica*, 60 (1): 79–84.

———. (2007). *Epistemology Modalized.* London: Routledge.

———. (2008). "Epistemic Luck and the Generality Problem." *Philosophical Studies*, 139

(3): 353–366.

———. (2009). "Margins for Error and Sensitivity: What Nozick Might Have Said." *Acta Analytica*, 24 (1): 17–31.

———. (2012a). "Basic Knowledge and Easy Understanding." *Acta Analytica*, 27: 145–161.

———. (2012b). "Methods and How to Individuate Them." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 81–97.

———. (2013). "Why Reliabilism Does Not Permit Easy Knowledge." *Synthese*, 190 (17): 3751–3775.

———. (2016). "BIVs, Sensitivity, Discrimination, and Relevant Alternatives." In Sanford C. Goldberg (ed.), *The Brain in a Vat*. New York: Cambridge University Press, 111–127.

———. (2018). "The Sensitivity Response to the Gettier Problem." In Stephen Hetherington (ed.), *The Gettier Problem*. Cambridge: Cambridge University Press, 108–124.

Beddor, Bob, & Pavese, Carlotta. (2020). "Modal Virtue Epistemology." *Philosophy and Phenomenological Research*, 101 (1): 61–79.

Bernecker, Sven. (2012). "Sensitivity, Safety, and Closure." *Acta Analytica*, 27 (4): 367–381.

———. (2020). "Against Global Method Safety." *Synthese*, 197 (12): 5101–5116.

Bjerring, Jens Christian, & Gundersen, Lars Bo. (2020). "Higher-Order Knowledge and Sensitivity." *Canadian Journal of Philosophy*, 50 (3): 339–349.

Black, Tim. (2002). "A Moorean Response to Brain-in-a-Vat Scepticism." *Australasian Journal of Philosophy*, 80 (2): 148–163.

———. (2008). "Defending a Sensitive Neo-Moorean Invariantism." In Vincent F. Hendricks

& Duncan Pritchard (eds.), *New Waves in Epistemology*. London: Palgrave Macmillan, 8–27.

———. (2011). "Modal and Anti-Luck Epistemology." In Sven Bernecker & Duncan Pritchard (eds.), *The Routledge Companion to Epistemology*. New York: Routledge, 187–198.

———. (2019). "Anti-Luck Epistemology." In Ian M. Church & Robert J. Hartman (eds.), *The Routledge Handbook of the Philosophy and Psychology of Luck*. New York: Routledge, 284–294.

Black, Tim, & Murphy, Peter. (2007). "In Defense of Sensitivity." *Synthese*, 154 (1): 53–71.

Blome-Tillmann, Michael. (2017). "Sensitivity Actually." *Philosophy and Phenomenological Research*, 94 (3): 606–625.

Bogardus, Tomas, & Marxen, Chad. (2014). "Yes, Safety is in Danger." *Philosophia*, 42 (2): 1–14.

Bolos, Anthony, & Collin, James Henry. (2018). "A Sensitive Virtue Epistemology." *Synthese*, 195 (3): 1321–1335.

Broncano-Berrocal, Fernando. (2014). "Is Safety In Danger?" *Philosophia*, 42 (1): 1–19.

———. (2018). "Knowledge and Tracking Revisited." *Analysis*, 78 (3): 396–405.

Carter, Joseph Adam. (2010). "Anti-Luck Epistemology and Safety's Discontents." Philosophia, 38 (3): 517–532.

Collin, James Henry. (2018). "Towards an Account of Epistemic Luck for Necessary Truths." *Acta Analytica*, 33 (4): 483–504.

Collin, James Henry, & Bolos, Anthony. (2020). "Sensitivity Theorists Aren't Unhinged."

Erkenntnis, https://doi.org/10.1007/s10670-019-00206-2

Cross, Troy. (2010). "Skeptical Success." In Tamar Szabó Gendler & John Hawthorne (eds.) *Oxford Studies in Epistemology: Volume 3*. Oxford: Oxford University Press, 35–62.

David, Marian, & Warfield, Ted A. (2008). "Knowledge-Closure and Skepticism." In Quentin Smith (ed.), *Epistemology: New Essays*. Oxford: Oxford University Press, 137–187.

DeRose, Keith. (1995). "Solving the Skeptical Problem." *Philosophical Review*, 104 (1): 1–52.

———. (2010). "Insensitivity Is Back, Baby!" *Philosophical Perspectives*, 24: 161–187.

Dutant, Julien. (2010). "Two Notions Of Safety." *Swiss Philosophical Preprints*, 87: 1–19.

———. (2016). "How to be an Infallibilist." *Philosophical Issues*, 26 (1):148–171.

Garrett, Brian J. (1983). "Nozick on Knowledge." *Analysis*, 43 (3): 181–184.

Goldberg, Sanford. (2012). "Sensitivity from Others." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 43–65.

Goldman, Alvin. (1976). "Discrimination and Perceptual Knowledge." *Journal of Philosophy*, 73 (20): 771–791.

———. (2009). "Williamson on Knowledge and Evidence." In Patrick Greenough & Duncan Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press, 73–91.

Greco, John. (2012). "Better Safe than Sensitive." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press, 193–206.

Grundmann, Thomas. (2020). "Saving Safety From Counterexamples." *Synthese*, 197:

5161–5185.

Gundersen, Lars Bo. (2003). *Dispositional Theories of Knowledge*. Burlington: Ashgate.

———. (2010). "Tracking, Epistemic Dispositions and the Conditional Analysis." *Erkenntnis*, 72: 353–364.

———. (2012). "Knowledge, Cognitive Dispositions and Conditionals." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 66–81.

Hawthorne, John. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hirvelä, Jaakko. (2017). "Is it Safe to Disagree?" *Ratio*, 30: 305–321.

———. (2019). "Global Safety: How to Deal with Necessary Truths." *Synthese*, 196 (3): 1167–1186.

———. (2020). "No Safe Haven for the Virtuous." *Episteme*, 17 (1): 48–63.

Huemer, Michael. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.

Ichikawa, Jonathan Jenkins. (2011). "Quantifiers, Knowledge, and Counterfactuals." *Philosophy and Phenomenological Research*, 82: 287–313.

Kelp, Christoph. (2013). "Knowledge: The Safe-Apt View." *Australasian Journal of Philosophy*, 91 (2): 265–278.

Kripke, Saul. (2011). "Nozick on Knowledge." In *Philosophical Troubles: Collected Papers, Volume 1*. New York: Oxford University Press, 162–224.

Kvanvig, Jonathan L. (2006). "Closure Principles." *Philosophy Compass*, 1 (3): 256–267.

Lasonen-Aarnio, Maria. (2010). "Unreasonable Knowledge." *Philosophical Perspectives*, 24

(1): 1–21.

Lewis, David. (1973a). *Counterfactuals*. Oxford: Blackwell.

———. (1973b). "Counterfactuals and Comparative Possibility." *Journal of Philosophical Logic*, 2 (4): 418–446.

Luper, Steven. (2003). "Indiscernability Skepticism." In Steven Luper (ed.), *The Skeptics: Contemporary Essays (Ashgate Epistemology and Mind Series)*. Burlington: Ashgate Publishing, 183–202.

———. (2006a). "Dretske on Knowledge Closure." *Australasian Journal of Philosophy*, 84 (3): 379–394.

———. (2006b). "Restorative Rigging and the Safe Indication Account." *Synthese*, 153 (1): 161–170.

———. (2012). "False Negatives." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 207–226.

———. (2020). "Epistemic Closure." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, URL = <https://plato.stanford.edu/archives/sum2020/entries/closure-epistemic/>

Luper-Foy, Steven. (1984). "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism." *Australasian Journal of Philosophy*, 62 (1): 26–49.

Manley, David. (2007). "Safety, Content, Apriority, Self-Knowledge." *Journal of Philosophy*, 104 (8): 403–423.

Melchior, Guido. (2015). "The Heterogeneity Problem for Sensitivity Accounts." *Episteme*, 12 (4): 479–496.

———. (2017a). "Epistemic Luck and Logical Necessities: Armchair Luck Revisited." In Bojan Borstner & Smiljana Gartner (eds.), *Thought Experiments between Nature and Society: A Festschrift for Nenad Miščević*. Cambridge: Cambridge Scholars Publishing, 137–150.

———. (2017b). "Sensitivity Has Multiple Heterogeneity Problems: A Reply to Wallbridge." *Philosophia*, 45 (4): 1741–1747.

———. (2021). "Sensitivity, Safety, and Impossible Worlds." *Philosophical Studies*, 178: 713–729.

Miščević, Nenad. (2007). "Armchair Luck: Apriority, Intellection and Epistemic Luck." *Acta Analytica*, 22 (1): 48–73.

Murphy, Peter, & Black, Tim. (2012). "Sensitivity Meets Explanation: An Improved Counterfactual Condition on Knowledge." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 28–42.

Neil, Charles. (2021). "Safety, Domination, and Differential Support." *Synthese*, 198: 1139–1152.

Nozick, Robert. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.

Paterson, Niall J. (2020). "Safety and Necessity." *Erkenntnis*. https://doi.org/10.1007/s10660-020-00231-6

Peet, Andrew, & Pitcovski, Eli. (2018). "Normal Knowledge: Toward an Explanation-Based Theory of Knowledge." *Journal of Philosophy*, 115 (3): 141–157.

Pritchard, Duncan. (2002). "Resurrecting the Moorean Response to Scepticism." *International Journal of Philosophical Studies*, 10: 283–307.

———. (2005). *Epistemic Luck*. Oxford: Oxford University Press.

———. (2007). "Anti-Luck Epistemology." *Synthese*, 158: 277–298.

———. (2008a). "Knowledge, Luck, and Lotteries." In Vincent F. Hendricks & Duncan Pritchard (eds.). *New Waves in Epistemology*. London: Palgrave Macmillan, 28–51.

———. (2008b). "Sensitivity, Safety, and Anti-Luck Epistemology." In John Greco (ed.), *Oxford Handbook of Skepticism*. Oxford: Oxford University Press, 437–455.

———. (2008c). "Virtue Epistemology and Epistemic Luck, Revisited." *Metaphilosophy*, 39 (1): 66–88.

———. (2009). "Safety-Based Epistemology: Whither Now?" *Journal of Philosophical Research*, 34: 33–45.

———. (2012a). "Anti-Luck Virtue Epistemology." *Journal of Philosophy*, 109 (3): 247–279.

———. (2012b). "In Defence of Modest Anti-Luck Epistemology." In Tim Black & Kelly Becker (eds.), *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press, 173–192.

———. (2013). "Knowledge Cannot Be Lucky." In Matthias Steup, John Turri, & Ernest Sosa (eds), *Contemporary Debates in Epistemology, 2nd edition*. Malden, MA: Blackwell, 152–164.

———. (2015). "Anti-Luck Epistemology and the Gettier Problem." *Philosophical Studies*, 172: 93–111.

———. (2016). *Epistemology*, *2nd edition*. New York: Palgrave Macmillan.

———. (2017). "Knowledge, Luck, and Virtue: Resolving the Gettier Problem." In Rodrigo

Borges, Claudio de Almeida, & Peter Klein (ed.), *Explaining Knowledge: New Essays on The Gettier Problem*. Oxford: Oxford University Press, 57–73.

———. (2018). "The Gettier Problem and Epistemic Luck." In Stephen Hetherington (ed.), *The Gettier Problem*. Cambridge: Cambridge University Press, 96–107.

Ramachandran, Murali. (2015). "Knowing by Way of Tracking and Epistemic Closure." *Analysis*, 75 (2): 217–223.

Roland, Jeffrey, & Cogburn, Jon. (2011). "Anti-Luck Epistemologies and Necessary Truths." *Philosophia*, 39 (3): 547–561.

Roush, Sherrilyn. (2005). *Tracking Truth: Knowledge, Evidence, and Science*. Oxford: Clarendon Press.

———. (2012). "Sensitivity and Closure." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*. New York: Cambridge University Press, 242–268.

Sainsbury, Richard M. (1997). "Easy Possibilities." *Philosophy and Phenomenological Research*, 57 (4): 907–919.

Schulz, Moritz. (2020). "Finding Closure for Safety." *Episteme*, 1–15.

Sosa, Ernest. (1996). "Postscript to 'Proper Functionalism and Virtue Epistemology.' " In Jonathan Kvanvig (ed.), *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledg*e. Lanham: Rowman & Littlefield, 269–278.

———. (1999a). "How Must Knowledge Be Modally Related to What Is Known?" *Philosophical Topics*, 26 (1&2): 373–384.

———. (1999b). "How to Defeat Opposition to Moore." *Philosophical Perspectives*, 13: 141–154.

———. (2002). "Tracking, Competence, and Knowledge." In Paul K. Moser (ed.) *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press, 264–287.

———. (2003). "Neither Contextualism Nor Skepticism." In Steven Luper (ed.), *The Skeptics: Contemporary Essays (Ashgate Epistemology and Mind Series)*. Burlington: Ashgate Publishing, 165–182.

———. (2004). "Relevant Alternatives, Contextualism Included." *Philosophical Studies*, 119: 35–65.

———. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.

———. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford: Oxford University Press.

———. (2015). *Judgment and Agency*. Oxford: Oxford University Press.

Stalnaker, Robert C. (1968). "A Theory of Conditionals." In Nicolas Rescher (ed.), *Studies in Logical Theory*. Oxford: Blackwell, 98–112.

Topey, Brett. (2021). "Saving Sensitivity." *Philosophical Quarterly,* https://doi.org/10.1093/pq/pqab015

Vogel, Jonathan. (1987). "Tracking, Closure, and Inductive Knowledge." In Steven Luper (ed.), *The Possibility of Knowledge: Nozick and His Critics*. Lanham: Rowman & Littlefield, 197–215.

———. (2000). "Reliabilism Leveled." *Journal of Philosophy*, 97 (11): 602–623.

———. (2007). "Subjunctivitis." *Philosophical Studies*, 134 (1): 73–88.

———. (2012). "The Enduring Trouble with Tracking." In Kelly Becker & Tim Black (eds.),

The Sensitivity Principle in Epistemology. Cambridge: Cambridge University Press, 122–151.

Wallbridge, Kevin. (2017). "Sensitivity Hasn't Got a Heterogeneity Problem: A Reply to Melchior." *Philosophia*, 45 (2): 835–841.

———. (2018a). "Sensitivity and Higher-Order Knowledge." *Pacific Philosophical Quarterly*, 99 (2): 164–173.

———. (2018b). "Sensitivity, Induction, and Miracles." *Australasian Journal of Philosophy*, 96 (1): 118–126.

Warfield, Ted A. (2004) "When Epistemic Closure Does and Does Not Fail: A Lesson from the History of Epistemology." *Analysis*, 64 (1): 35−41.

Wedgwood, Ralph. (2020). "The Internalist Virtue Theory of Knowledge." *Synthese*, 197: 5357–5378.

Williams, Michael. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.

Williamson, Timothy. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

———. (2009). "Reply to John Hawthorne and Maria Lasonen-Aarnio." In Patrick Greenough & Duncan Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press, 313–329.

Zalabardo, José. (2012). *Scepticism and Reliable Belief*. Oxford: Oxford University Press.

———. (2017). "Safety, Sensitivity and Differential Support." *Synthese*, 197 (12): 5379–5388.

Zhao, Bin. (2021a). "A Dilemma for Globalized Safety." *Acta Analytica*.

https://doi.org/10.1007/s12136-021-00478-w

———. (2021b). "Knowledge from Falsehood, Ignorance of Necessary Truths, and Safety." *Philosophia*. https://doi.org/10.1007/s11406-021-00410-x

———. (forthcoming). "Epistemic Closure, Necessary Truths, and Safety." *American Philosophical Quarterly*.

Zhao, Haicheng. (2020). "Knowledge Without Safety." *Synthese*, 197: 3261–3278.

———. (2021). "Better Virtuous than Safe." *Synthese*, 198: 6969–6991.

Zhao, Haicheng, & Baumann, Peter (2021). "Inductive Knowledge and Lotteries: Could One Explain Both 'Safely'?" *Ratio*, 34 (2): 118–126.