Synthese (2019) 196:1167–1186
https://doi.org/10.1007/s11229-017-1511-z

Global safety: how to deal with necessary truths

Jaakko Hirvelä[1]

Received: 3 February 2017 / Accepted: 24 July 2017 / Published online: 1 August 2017

Abstract According to the safety condition, a subject knows that p only if she would believe that p only if p was true. The safety condition has been a very popular necessary condition for knowledge of late. However, it is well documented that the safety condition is trivially satisfied in cases where the subject believes in a necessary truth. This is for the simple reason that a necessary truth is true in all possible worlds, and therefore it is true in all possible worlds where it is believed. But clearly, all beliefs concerning necessary truths do not amount to knowledge. The safety theorists have attempted to deal with the problem caused by necessary truths by restricting the safety condition to purely contingent truths and by globalizing the safety condition to a set of propositions. Both of these solutions are problematic. The principal aim of this paper is to develop a version of the safety condition that is able to deal with cases featuring necessary truths.

Keywords Safety · Necessary truths · Virtue epistemology · Generality problem · D. Pritchard · T. Williamson

1 Introduction

The idea behind the safety condition is that in order to know, one must be safe from error. Not only must one get things right as a matter of fact, it also has to be the case that one could not easily have got things wrong either. The safety condition can be formulated as a subjunctive conditional, which says that S's belief that p is safe if,

✉ Jaakko Hirvelä
jaakko.hirvela@helsinki.fi

[1] Department of Philosophy, History, Culture and Art Studies, University of Helsinki, Unioninkatu 38, PL 24, 00014 Helsinki, Finland

and only if, "S would believe that p only if it were so that p" (Sosa 1999, p. 142). This characterization of the safety condition is usually relativized to methods of belief formation and fleshed out in terms of a possible worlds-heuristic as follows:

SAFETY: S's belief that p is safe if, and only if,

(i) in most nearby possible worlds, and in all of the very closest possible worlds, in which S believes that p, via the same method of belief formation M that S uses in the actual world, p is true.

SAFETY is a modal condition for knowledge that seeks to eliminate the kind of luck that is incompatible with knowledge. As an anti-luck condition it is directly supported by Pritchard's modal analysis of luck (Pritchard 2005). Furthermore, SAFETY delivers correct verdicts regarding a wide range of thought experiments. It is able to deal with Gettier cases, lottery cases and with more sophisticated Gettier cases, like the famous barn façade case. Nevertheless, there is a set of cases that the SAFETY condition is unable to deal with. In order to have a specific case in mind, let me introduce the following case:

MALFUNCTION: Paige is using a calculator to count the product of 12×13. After punching in the equation the screen reads $12 \times 13 = 156$. Paige forms the corresponding (necessarily) true belief. Unbeknownst to Paige, however, the calculator is malfunctioning and generating answers at random.

Intuitively, Paige does not know that $[12 \times 13 = 156]$. After all, one cannot gain knowledge by using a malfunctioning calculator. However, given that $[12 \times 13 = 156]$ is a necessary truth, there are no possible worlds where Paige believes that $[12 \times 13 = 156]$ while her belief is false. Thus, SAFETY is trivially satisfied in cases where the subject believes in a necessary truth.[1]

While the problem of necessary truths does not threaten the necessity of SAFETY, it does undermine the sufficiency of the condition, and hence knowledge cannot be defined simply as a safe belief. However, even those epistemologists who think that SAFETY is merely a necessary condition for knowledge, have an interest in reformulating SAFETY in such a manner that it is not trivially satisfied if the subject believes in a necessary truth. The reason for this is that SAFETY is often put forward as an anti-luck condition for knowledge, that captures the sense in which knowledge is incompatible with luck, and it seems that it is a matter of luck that Paige believes the truth in MALFUNCTION. Therefore, in virtue of providing the necessary and sufficient conditions for avoidance of luck, SAFETY should be able to deal with this case as well.

The proponents of SAFETY have attempted to deal with the problem caused by necessary truths in two ways. Pritchard (2005) early attempt to deal with the problem was to restrict SAFETY to purely contingent truths. Later Pritchard (2012) has

[1] Roland and Cogburn (2011) have argued that safety conditions are unable to deal with cases featuring necessary truths.

abandoned this restricted version and demanded that in order for the subject to be properly safe from error, the belief formation method that she uses in the actual world cannot produce as an output any false beliefs in the scope of nearby possible worlds. Williamson (2009b) agrees that SAFETY should be globalized to a set of propositions in which the subject could easily have formed a belief in nearby possible worlds via the method of belief formation that she uses in the actual world. However, Williamson's formulation of SAFETY differs from Pritchard's, in that Williamson restricts the relevant set of propositions by demanding that all of the relevant propositions have to be close to each other. I will argue that the solutions offered by Pritchard and Williamson are unsatisfactory and that a better solution is available.

In Sect. 2 I will critically evaluate the prospects of restricting SAFETY to contingent propositions. In Sect. 3 I will present the way in which Pritchard seeks to globalize SAFETY. I will argue that the way in which Pritchard attempts to restrict the relevant set of propositions hinges on a wrongheaded account of how to individuate methods of belief formation. In Sect. 4 I will review Williamson's solution. I will argue that Williamson is correct in that the relevant set of propositions cannot satisfactorily be restricted solely in terms of the method of belief formation employed by the subject in the actual world, but that he has stumbled in the details of his positive view. Finally, in Sect. 5 I present my positive proposal. I will argue that safety theorists should relativize SAFETY to virtuous methods of belief formation that the subject uses in the actual world. Relativizing SAFETY to virtuous methods of belief formation allows the safety theorist to solve the generality problem, but it does not restrict the relevant set of propositions robustly enough. We need to restrict the relevant set of propositions through other means as well. I will argue that the relevant set of propositions should be restricted in terms of the agent's subject matter of inquiry. In order to be properly safe, it must be the case that one could not have easily ended up with a false belief in one's inquiry.

2 Contingentist solution

In his early work, Pritchard tried to evade the problem caused by necessary truths by confining his theory of knowledge to contingent propositions:

> I am only offering an elucidation of knowledge as it applies to fully contingent propositions, and so the most obvious counterexamples—those that, for example, appeal to nomic necessities and stubborn, but groundless, beliefs in these necessities in order to ensure that the target beliefs are safe—will not be applicable here. (2007, p. 290)

Restricting one's analysis of knowledge to contingent propositions in this way raises two questions: "is there a principled reason for restricting one's account to contingent propositions" and "is SAFETY able to deal with all cases featuring contingent propositions"? I will argue that there is no principled reason for restricting one's account of knowledge to contingent propositions and hence Pritchard's move is *ad hoc*. Furthermore, I will demonstrate that SAFETY is unable to deal with some cases featuring beliefs in contingent truths for the same reason that it is unable to deal with necessary

truths. Restricting SAFETY to contingent truths does not help the safety theorist at all. Let us tackle the second issue first.

Many contingently true propositions are true in all nearby possible worlds, but false in some faraway possible worlds. My belief that I am not a brain in a vat is true in all nearby possible worlds because, given the current technology, it is impossible to produce brains in vats that have experiences that are as vivid as my current experiences. Therefore, there are no nearby possible worlds where I am a brain in a vat, and hence my belief that I am not, satisfies SAFETY, regardless of the way in which I arrived at my belief. But clearly there are some deviant ways of forming this belief that would not give rise to knowledge. What if the belief was a result of wishful thinking?

If one is inclined to think that we always know the denial of sceptical hypotheses, regardless of the way in which the belief was acquired, one might not be persuaded by the above example. Greco (2007, pp. 301–302) has offered a non-sceptical scenario that will suit our purposes.

FROG: Fred is colourblind. Currently Fred is looking at a frog and forms the true belief that the frog is green. As it happens, all the frogs in Fred's environment are bound to be green, due to their evolutionary history. Moreover, frogs from other regions could not easily have been in Fred's environment. Therefore the frog that Fred is looking at is green in all nearby possible worlds, and Fred's belief that it is green is true in all nearby possible worlds where he holds onto that belief. Fred satisfies SAFETY, but clearly Fred does not know that the frog is green, because he is unable to distinguish between different colours.

Given that it is a contingent fact that the frog is green, SAFETY delivers wrong verdicts even in cases featuring contingently true propositions. Moreover, the reason why SAFETY is satisfied in FROG is the same reason why it is satisfied in MALFUNCTION. In both cases the subjects form their beliefs in a haphazard way, but due to the nature of the proposition believed, the subject's belief continues to be true across nearby possible worlds. When we assess whether a subject's belief is safe or not we are interested only in the nearby possible worlds. It does not matter that Paige's belief continues to be true in all possible worlds, whereas Fred's belief continues to be true only in nearby possible worlds. Since SAFETY is satisfied in FROG for the same reason that it is satisfied in MALFUNCTION restricting SAFETY to contingently true propositions is not a remedy.

But is there still a principled reason for restricting one's account of knowledge to contingent propositions? To see why Pritchard might have thought that the restriction is principled, it is important to note that Pritchard was aiming to formulate an anti-luck condition. Pritchard's anti-luck epistemology had the following agenda: first we need to give a general account of what luck is, then we need to consider what kind of luck is incompatible with knowledge, and after that we can formulate an anti-luck condition that ensures that no luckily true belief will amount to knowledge. Pritchard analyzed lucky events as follows: it is a matter of luck for subject S that event E occurred only if E is significant to S, and E occurs in the actual world but fails to occur in most nearby possible worlds where the relevant initial conditions for E are the same

(Pritchard 2005, pp. 129, 132).[2] The difference between good and bad luck resides in whether E is of positive or negative significance to S. Of course, events that are bound to occur will not be lucky in any sense on the modal account of luck. This analysis of lucky events, led Pritchard to advocate SAFETY as an adequate anti-luck condition for knowledge.

However, it is reasonable to doubt whether Pritchard managed to capture what it means for an event to be lucky for someone. Suppose that a game between two football teams ended 2–0. Suppose further that in most nearby possible worlds the game did not end 2–0, and that the fact that the game ended 2–0 is of positive significance to the fans of the winning team. In Pritchard's analysis it should be a matter of good luck for the fans of the winning team that the game ended 2–0. But assume that in most nearby possible worlds the team that won would have won the game even more clearly. Is it then a matter of luck for the fans that the game ended 2–0? I suggest not. The reason why it is not a lucky event for the fans is that similar events, equal or greater in significance, would have occurred if the actual event had not occurred. Coffman agrees. He writes that in order for an event E to be lucky for a subject, it must be the case that "there was just before t a large chance that no event sufficiently similar and equal in significance to E would occur at t" (2007, p. 396).

If Coffman is right about this, then when evaluating whether an event E is lucky for S we should assess not only the modal profile of E, but also the modal profile of events that are similar to E and at least equally significant. The same should apply when we evaluate whether it is a matter of luck that the subject has a true belief. We should not confine ourselves to assessing only whether the belief that the subject formed in the actual world remains true across possible worlds where it is believed; rather, we should also take into account the possible different beliefs that the subject could have arrived at in her inquiry, and consider, whether those beliefs also continue to be true in possible worlds where the subject has those beliefs. But to make this move we must globalize SAFETY to a set of propositions. What is important to note, is that this move can be motivated solely by reflecting on the nature of lucky events. In the next section I will consider the way in which Pritchard seeks to globalize the safety condition.

3 Pritchard's globalization

In his more recent work Pritchard has not restricted his analysis of knowledge to contingently true propositions.[3] Instead, he has argued that when evaluating whether a subject's belief is safe or not, we must consider whether the subject could easily have arrived at a false belief by forming her belief on the same basis on which she formed her belief in the actual world. Pritchard writes that "For what we are interested in is rather how the agent forms her beliefs in similar circumstances and in response to the

[2] In his later work Pritchard (2015) has rejected the significance condition.

[3] Pritchard (2012) no longer thinks that knowledge can be analysed solely in terms of the safety condition but has added a separate virtue-theoretic condition to his analysis.

same stimuli. These beliefs may be beliefs *that p*, but equally they may be beliefs in distinct propositions" (2012, pp. 256–257).

Accordingly, the safety condition that Pritchard currently endorses can be formulated as follows:

SAFETY^P: S's belief that p, formed on basis B, is safe, if and only if:

(i) in most nearby possible worlds, and in all of the very closest possible worlds, where S forms a belief on basis B, S's belief is true.

It should be easy enough to see how this qualification enables Pritchard to deal with MALFUNCTION and FROG. In both cases, the subjects could easily have believed in a different proposition, which would have been false, on the same basis. Fred could easily have believed that something that is red is green by looking, while Paige could easily have believed that [12 × 13 = 146] by using the malfunctioning calculator. However, it is fair to ask how the basis B is to be individuated. Obviously, the extension of SAFETY^P will vary greatly between fine and coarse-grained individuations of the basis of belief formation. Individuating basis in a coarse-grained way runs the risk of scepticism since the subject will have to have a safe epistemic stance regarding a very wide range of propositions, while individuating the basis in a fine-grained way threatens to make the globalization of SAFETY futile. In order to see whether SAFETY^P is plausible, we need to consider the way in which Pritchard individuates the basis of belief formation.

Sadly, Pritchard has not said much about this issue. His remarks on how to individuate the basis of belief formation are remarkably vague. What is clear is that he seeks to individuate the basis in an externalist fashion, since he writes that "the 'way' in which the belief is actually formed needs to be individuated *externally* rather than *internally*" (2005, p. 152). An individuation of the basis of belief formation is external if the basis can include factors that are completely external to the subject's phenomenal states, and internal if it can include only factors that are internal to the subject's phenomenal states. An external individuation of basis can include the neural process that is responsible for the belief, the stimuli that caused the belief (photons hitting the retina), or the object responsible for the stimuli that caused the belief, etc. While Pritchard does not commit himself to any explicit way of individuating the basis of belief formation, some clues can be gathered from the way in which he handles some thought experiments.

Commenting on the famous barn façade case, he writes that "had it been false that the object that Barney is looking at is a barn—that is, if it were a barn façade—Barney would have continued to believe this proposition regardless, since his basis for this belief (that he is presented with a plausible barn-shaped object) would be unchanged" (Pritchard 2012, p. 251).[4] This passage suggests that in cases of perceptual beliefs, the basis that is kept fixed is the fact that one is presented with a similar looking object that one saw in the actual world. This way of individuating the basis delivers the correct

[4] Note that "being presented with a barn-shaped object" can be interpreted along internalist lines, as a barn-shaped being present to one's consciousness, or along externalist lines as a barn-shaped object being present in one's (physical) environment. Given Pritchard's claim that the basis should be understood externally I take it that he has the later interpretation in mind.

verdict regarding the barn façade case, but runs into problems with the following kinds of cases:

BOWL OF FRUIT: Sofia is looking at what seems to be a bowl of fruit and forms the true belief that [that is an apple] while pointing at an apple. Sofia has excellent eyesight and the perceptual conditions are fine. However, as it happens, only the apples in the bowl are real. The other 'fruits' in the bowl are very realistic plastic replicas that would have fooled anyone. Luckily, Sofia happened to point at an apple.

Clearly, Sofia does not have knowledge in this case. After all, she could very easily have gained a false belief, and therefore it is a matter of luck that her belief is true. However, if a part of Sofia's basis of belief formation is that she is presented with an apple-like object, then in all nearby possible worlds where the basis of her belief is kept fixed her belief will continue to be true. Therefore she satisfies SAFETY^P. This way of individuating the basis restricts the relevant set of propositions too severely.

Commenting on another thought experiment introduced by Bogardus (2014, pp. 300–301), where a subject forms a belief about the time by consulting the world's most reliable clock, which could have paradoxically been malfunctioning and telling the wrong time, Pritchard writes:

> The crux of the matter is that we need to keep the subject's actual evidential basis fixed, and of course his actual evidential basis for his belief is formed by consulting the reliable and unaffected clock. While it is lucky that the subject has this evidential basis (in that there are close possible worlds where it is absent), it is not lucky that she forms a true belief on this basis. Indeed, in all close possible worlds where she continues to enjoy the same evidential basis she continues to form a true belief. (Pritchard 2016, p. 9)

This passage seems to suggest even stronger requirements on the individuation of the basis and highlights the fact that Pritchard individuates methods externally, rather than internally. Not only must the subject form her belief by consulting the same clock in the relevant counterfactual scenarios but the clock must also function properly in the relevant possible worlds. Clearly, it will not be enough that the subject is presented with a similar looking object, since the clock will look just the same whether it is functioning or not. Allowing the functional integrity of an external device to be a part of the basis of belief formation can, perhaps, be motivated by the extended mind hypothesis,[5] but such a move, combined with SAFETY^P will deliver wrong verdicts regarding a wide range of cases. BOWL OF FRUIT will still be a counterexample, and there are others as well. Consider the following case adapted from Goldman (1986, p. 45):

THERMOMETERS: Tom is feeling feverish and decides to check his body temperature. In actual fact, Tom has a body temperature of 38.2 °C. Tom has a box full of thermometers which he justifiably takes to be reliable. Tom picks one of them at random, consults it, and forms the true belief that his body temperature is 38.2 °C. The twist in the tale, however, is that the thermometer that Tom happened to choose

[5] See Clark and Chalmers (1998) for the extended mind hypothesis.

was the only functioning thermometer in the box. All of the others show relatively plausible temperatures, which are, nevertheless, false.

Clearly, Tom does not know his body temperature in this case. He could very easily have picked a malfunctioning thermometer and gained a false belief instead. The possible world where he picks a malfunctioning thermometer ought to be a relevant possible world when evaluating the safety of his belief. However, if we individuate basis along the lines suggested by Pritchard above, Tom's belief will end up being safe. After all, there are no nearby possible worlds where Tom gains a false belief by using the properly functioning thermometer that he uses in the actual world. Therefore, Pritchard's way of individuating the basis delivers a wrong verdict regarding this case as well.[6]

Perhaps there is a way of individuating methods of belief formation that will allow the safety theorist to restrict the relevant set of propositions in a satisfactory manner, but it is not the one that can be extracted from the remarks offered by Pritchard. In the next section, we will examine how Williamson aims to restrict the relevant set of propositions.

4 Williamson's globalization

Williamson offers different formulations of the safety principle at various places. Here is the most promising formulation:

SAFETY^W: S's belief that p, formed on basis B, is safe in case α if and only if,

(i) S believes p on basis B in α, and in any case close to α in which S believes a proposition $p*$ close to p on basis close to B, $p*$ is true. (Williamson 2009b, p. 325).[7]

At first glance, SAFETY^W might be able to overcome the problems that SAFETY^P succumbed to. Since Williamson restricts the relevant set of propositions, not only in terms of the basis of belief formation that the subject used in the actual world, but also in terms of 'propositional closeness', and he allows for small variance in the basis of belief formation, he can afford to individuate the basis in a more coarse-grained fashion than Pritchard could. Moreover, it is important to note that Williamson, unlike Pritchard, is not trying to provide an analysis of knowledge in terms of safety. Williamson does not think that we can always (or even often) determine whether a belief is safe before considering whether it amounts to knowledge. This is a consequence of his knowledge first epistemology. In what follows, we will examine the merits of SAFETY^W as a

[6] Perhaps Pritchard could claim that BOWL OF FRUIT and THERMOMETERS are actually cases of knowledge since they feature only evidential epistemic luck, which is compatible with knowledge possession. By executing such a maneuver, Pritchard would concede that his theory gives unintuitive verdicts regarding certain cases. However, there is reason to doubt that Pritchard would be inclined to make such a move, since he thinks that the barn façade case features environmental luck, which is knowledge undermining, and the BOWL OF FRUIT case is very similar to the barn façade case. Furthermore, even if THERMOMETERS did feature evidential luck, we cannot infer from this fact alone that it does not feature a knowledge undermining luck also.

[7] A 'case' is a subject-centred possible world. Here after I will refer to cases as possible worlds for the sake of simplification.

necessary condition for knowledge, but one that is possibly circular. Moreover we should bear in mind that safety conditions are supposed to be able to rule out the kind of luck that is harmful for knowledge, and it is therefore reasonable to think that SAFETY^W would be able to deal with cases like MALFUNCTION.

In order to determine whether SAFETY^W restricts the relevant set of propositions in a satisfactory manner, we need to answer two questions. "How is basis B individuated?" and "What does it mean for a proposition to be close to another one?" Let us look at the initial question first.

Given that Williamson does not require that the basis has to be the same in all relevant possible worlds, but that it has to be 'close' to the basis that the subject had in the actual world, one might think that the conception of basis that Williamson deploys is much wider than the one advocated by Pritchard. This impression is false. To see this, consider Williamson's reply to the following case of unsafe knowledge offered by Neta and Rohrbaugh (2004, pp. 399–400):

> WATER: I am drinking a glass of water which I have just poured from the bottle. Standing next to me is a happy person who has just won the lottery. Had this person lost the lottery, she would have maliciously polluted my water with a tasteless, odorless, colorless toxin. But since she won the lottery, she does no such thing. Nonetheless, she almost lost the lottery. Now, I drink the pure, unadulterated water and judge, truly and knowingly, that I am drinking pure, unadulterated water. But the toxin would not have flavored the water, and so had the toxin gone in, I would still have believed falsely that I was drinking pure, unadulterated water.

Commenting on this case, and another very similar to it, Williamson writes:

> [The cases] involve large external differences between the processes that actually generate the beliefs in question in the good cases and the processes that generate errors in the corresponding bad cases. Again, these differences turn out to constitute large enough dissimilarities between the bases of one's belief to allow one to know: the cases are not similar enough to exclude safety. (2009a, p. 307)

But if cases like WATER involve too big differences in the basis in the good case where the subject drinks water, and in the bad case where the subject drinks the toxin, then surely the subject will form her belief on a different basis when she sees a real barn and when she sees the perceptually equivalent barn façade. But if SAFETY^W is to be able to deliver the correct verdict regarding the barn façade case, then the subject must have the same basis for her belief in the case where she sees a real barn and a façade. If Williamson individuates basis of belief formation along the above lines, SAFETY^W is unable to deal with a wide variety of cases that it should be able to deal with.

But perhaps Williamson can claim that in WATER there are large external differences between the processes that generate the beliefs in the good and bad cases, while maintaining that in the barn façade case the difference is not that big. Williamson is free to make such a move because he is not aiming to provide a reductive analysis of

knowledge.[8] Williamson claims that the decision whether the bases are similar enough should be made *post hoc*, after considering whether the cases involve knowledge. Needless to say, versions of the safety condition that rely on *post hoc* methodology are inferior in explanatory power to versions that do not rely on such methodology. If we can shed more light on what counts as a proper basis in a non-*post hoc* way, we should do it.

Leaving open how Williamson individuates the basis of belief formation, we can move on to the second question: What does it mean for a proposition to be close to another one? Williamson does little to clarify the concept of propositional closeness, but the kind of examples that he uses when he applies the concept of safety might give us a clue:

> Some hunters see a deer disappear behind a rock. They believe truly that it is behind the rock. To complete their kill, they must maintain a true belief about the location of the deer for several minutes. But since it is logically possible for the deer to be behind the rock at one moment and not at another, their present-tensed belief may be true at one moment and false at another. By standard criteria of individuation, a proposition cannot change its truth-value; the sentence 'The deer is behind the rock' expresses different propositions at different times. […] If the hunters know that the deer is behind the rock, they have the kind of sensitivity to its location that makes them more likely to have future true belief about its location than they are if they merely believe truly that it is behind the rock. (Williamson 2000, p. 101)

In this example, the 'close' propositions are expressed by using the same sentence, namely "The deer is behind the rock". The famous anti-luminosity argument (Williamson 2000, pp. 96–98) features similar 'close' propositions that are expressed by using the same sentence.[9] But if the proposition p is close to the proposition $p*$ just in case p and $p*$ can be expressed by using the same sentence, then SAFETY^W cannot handle the cases that led us to globalize SAFETY in the first place or it leads into scepticism. The reason for this is that a sentence can be used to express a necessary truth and a necessary falsehood only if the sentence includes highly context sensitive words, such as demonstratives. For example, the sentence "the sum of those numbers is 4" can express a necessarily true or false proposition depending on what "those" refer to. However, if we allow the relevant sentence to include demonstratives the scope of relevant propositions will be far wider than desired. Moreover, in cases

[8] Of course Williamson could provide independent reasons for thinking that the relevant differences are big in cases like WATER, while they are not big enough in the barn façade case. An anonymous referee objected that Williamson could claim that the cases differ not just with respect to knowledge but also in the causal structure of how the subject gained a true belief. According to the referee "in the bad case in WATER, what's in the glass is causally influenced by a mental state of a third party that simply isn't present in the good case; whereas in the barn case, whether one is looking at a barn or a façade doesn't depend causally on any such thing." However, there is reason to doubt whether such differences are epistemically relevant. Luper (2006, pp. 164–165), for instance, has argued that whether the fact that one is in a good case or bad case depends on someone else's mental states, or on non-agential features of the case, does not make a difference to the epistemic state of the subject.

[9] See Williamson (1994, pp. 231–234) for similar safety-style reasoning.

such as MALFUNCTION the belief that Paige forms is $12 \times 13 = 156$, not that the product of those numbers is 156. Finally, subjects whose language lacked context sensitive expressions would always know a necessary truth when they believed in one, but surely the fact that our language contains context sensitive expressions does not explain (even partially) why a belief in a necessary truth is not automatically knowledge. Therefore, we should not unpack propositional closeness in this way.

Williamson could also say that the proposition p is close to the proposition $p*$ just in case $p*$ is true in some world $w*$ that is close to a world w where p is true. This reading of propositional closeness would suffice for the anti-luminosity argument and would fit Williamson's example. However, it ought to be clear that this reading would not help Williamson in the case of necessary truths. After all, in order to deal with such cases Williamson would have to maintain that a necessary falsehood is close to some necessary truth, but on this reading they could never be close to each other, since necessarily false propositions are not true in any possible world, while necessarily true propositions are true in all possible worlds.

Finally, one might think that propositions p and $p*$ are close to each other just in case they contain terms that are *semantically plastic* and p's extension could easily have been that of $p*$'s. According to Hawthorne, a term is semantically plastic just in case "(a) slight differences in usage make for differences in semantic value and (b) we are insensitive to the ways in which difference in usage makes for a difference in semantic value" (2006, p. 186). Vague terms like 'thin' and 'bald' are the best candidates for semantically plastic terms. If I were to believe that [Tim is thin], and Tim were a border-line case of being thin, I would not know that [Tim is thin] even if my belief was true. This is because the extension of 'thin' might easily have been such that Tim would not have belonged in its extension. However, even if this way of understanding propositional closeness allows the safety theorist to deal with problems caused by vague predicates, it will not help the safety theorist with respect to the problem of necessary truths. Terms like '12', '13' '=' 'x' and '156' are anything but vague. They are not semantically plastic, and therefore the semantic value of $[12 \times 13 = 156]$ could not easily have been the semantic value of $[12 \times 13 = 146]$. Moreover, if they were semantically plastic, wholesale scepticism would ensue, since one could not know that $[12 \times 13 = 156]$ even by using a properly functioning calculator.

Since it seems extremely hard to unpack what Williamson means by saying that a proposition is close to another one, perhaps propositional closeness should be taken as a primitive notion. Indeed, this seems to be the kind of move that we should allow Williamson to make, since he is not in the business of trying to analyse knowledge in the first place. However, taking propositional closeness as a primitive notion will leave SAFETY^W a very vague concept. Surely it is better to operate with explicit concepts if possible. If we can render SAFETY^W more explicit we should.

In order to provide a formulation of SAFETY that is able to deal with the problem of necessary truths without succumbing to scepticism, we need to look how the safety theorist should individuate methods of belief formation. I will argue that the safety theorist should demand that the subject uses the same virtuous method of belief formation in the relevant possible worlds. It is plausible that only virtuously formed beliefs can amount to knowledge, and therefore restricting our attention only to virtuously formed beliefs will do us no harm. Furthermore, I will argue that this way of individu-

ating methods of belief formation, while able to solve the generality problem, will not restrict the set of propositions enough.[10] Therefore, in addition to demanding that all of the relevant beliefs arise from the same virtuous method of belief formation that the subject uses in the actual world, we have to add a second restriction, as Williamson does. I propose that we restrict the relevant set of beliefs in terms of subject matters of inquiry. When assessing whether a subject knows something, we are interested in whether the subject could easily have gained a false belief in her inquiry. This way of restricting the relevant set of propositions is more informative and precise than the one suggested by Williamson.

5 Global safety

We saw above that the way in which Pritchard and Williamson sought to individuate methods of belief formation was too fine-grained. They delivered wrong verdicts regarding a wide range of cases. What we need is a more coarse-grained criterion of individuation. Before honing in on my positive proposal, it is useful to consider what kind of solution the generality problem requires. Conee and Feldman (1998, p. 3) claim that in order to answer the generality problem the reliabilist must provide a formula that picks a unique process type in each case. Reliabilists have tried to meet this demand by providing conditions of when two token processes can be treated as belonging to the same process type. Aiming to provide such an answer is, I believe, misguided.

The reason why we need not offer a general criterion of when two methods tokens are of the same type in order to give a satisfactory answer to the generality problem is because we are trying to provide an analysis of knowledge. If only certain methods of belief formation are knowledge-conducive, then we need to identify only those methods of belief formation. Being able to identify those methods of belief formation does not entail that we are able to give general criteria of when two methods of belief formation are the same. Therefore, by restricting our attention to only those methods of belief formation that are potentially knowledge-conducive, we will have made the task of answering the generality problem much easier.

It is plausible to think that only virtuously formed beliefs are in the market of knowledge. Virtue epistemologists, such as Sosa (2007), Greco (2010), Riggs (2007) and Zagzebski (1996), claim that knowledge is a cognitive achievement that is attributable to the exercise of the subject's cognitive abilities. Therefore, virtue epistemologists endorse the weaker claim that only virtuously formed beliefs can amount to knowledge. Let us assume that this weaker thesis is true, and that only beliefs that are gained through the exercise of epistemic competences or virtues can amount to knowledge.[11]

[10] The version of the generality problem that inflicts SAFETY is concerned with the question how the method of belief formation that the subject uses in the actual world should be individuated. Since the extension of safe beliefs will vary greatly between different ways of individuating methods of belief formation, as we have seen, a SAFETY-based theory of knowledge is radically incomplete without a solution to the generality problem.

[11] In what follows I will use the terms "competence" and "virtue" interchangeably.

Following Sosa (1991, p. 284), we can define cognitive virtue as a stable disposition D to attain true beliefs and avoid false belief within a certain field of propositions F, when in environment E and in conditions C. The field of propositions consists of propositions in which the relevant virtue could produce beliefs in. For example, the field of propositions in the case of good eyesight would consist of propositions such as [that is red], [that is square], etc. Dispositions, according to Sosa, have a three-part structure. They involve (i) constitution, (ii) condition, and (iii) situation. Cognitive virtues, or competences, as Sosa prefers to call them, follow suit. The constitution of a perceptual competence includes rods and cones; the condition component includes conditions such as being awake and clearheaded, and the situation includes conditions such as being in adequate lighting conditions. (Sosa 2010, pp. 465, 467)

Drawing on Sosa's virtue reliabilism, we can define when a subject's belief that p is virtuously formed as follows:

VIRTUE: A subject S's belief that p, which belongs to a field of propositions F, is virtuously formed via method V, in circumstance C if and only iff:

(i) S has an inner disposition D to attain correct doxastic attitudes with respect to propositions that belong to F,
(ii) S is in C and believes that p,
(iii) the fact that S believes that p, via V, is due to exercising D.

If VIRTUE captures what it takes for one's belief to be virtuously formed, the next question that we need to ask is "When are two virtuously formed beliefs the products of the same virtue?" Here we can draw on Broncano-Berrocal's work. Broncano-Berrocal (2014) has attempted to provide a solution to the generality problem on the behalf of the safety theorist. He argues that two methods m^1 and m^2 are the same only if they are globally reliable to the same extent, they are based on the same epistemic faculty, and that the circumstances in which the belief is gained through m^2 belong to the set of circumstances with respect to which m^1 is globally reliable (2014, p. 75).

Bogardus and Marxen (2014) have argued that the conditions proposed by Broncano-Berrocal are unsatisfactory for several reasons. The most serious objection that they raise is that Broncano-Berrocal's way of individuating methods of belief formation entails that there are no unreliable methods of belief formation. Moreover, they claim that in Broncano-Berrocal's view one cannot use the same method of belief formation in an environment which is suitable for the use of that method, and in an environment which is not suitable for that method.

In my view, the criticism that Bogardus and Marxen raise is apt. Clearly, there are many unreliable methods of belief formation. However, if we reformulate the conditions offered by Broncano-Berrocal not as conditions of when two methods of belief formation are the same, but as conditions of when two virtuously formed beliefs are products of the same virtuous method of belief formation, the critical points presented by Bogardus and Marxen turn into virtues. After all, there are no unreliable virtuous methods of belief formation. Furthermore, given that a belief is virtuously formed only in relation to the environment in which it is formed, a belief cannot be virtuously formed in an environment that is unsuitable for the use of that virtue.

Taking these points into account, we can individuate virtuously formed beliefs as follows:

METHOD: Two virtuously formed beliefs that p_1,and that p_2,which are the products of the virtuous methods of belief formation $v_1$and v_2, are the products of the same type of virtuous method V only if:

(i) v_1 and v_2 are globally reliable to the same degree with respect to the same field of propositions F (which includes p_1 and p_2), and the same range of circumstances C;
(ii) v_1 and v_2 are both based on the same epistemic faculty, such as vision or olfaction or audition or taction or gustation or testimony or deduction or induction or memory; and
(iii) the circumstances in which the belief p_2 is formed via v_2 are in the set of circumstances with respect to which v_1 is globally reliable and the circumstances in which the belief $p_1$is formed via $v_1$are in the set of circumstances with respect to which v_2 is globally reliable.

Assuming that only virtuously formed beliefs are in the market for knowledge, we can restrict SAFETY in terms of virtuous methods of belief formation as follows:

SAFETY^V: S's belief that p, formed via the virtuous method V, is safe if and only if:

(i) in most nearby possible worlds, and in all of the very closest possible worlds, where S forms a belief via V, S's belief is true.[12,13]

SAFETY^V yields correct verdicts regarding the cases that we have considered so far. In MALFUNCTION Paige uses a virtuous method of belief formation. Relying on a calculator is a globally reliable way of forming mathematical beliefs. True, the calculator that Paige is using is malfunctioning, but this does not undermine the global reliability of forming beliefs via calculators. Secondly, Paige is using the calculator in circumstances that are suitable for the use of calculators. Thirdly, Paige believes that [12 × 13 = 156] because she exercises her inner disposition to rely on the calculator. Therefore, Paige's belief is virtuously formed in the actual world. Moreover, she uses the same virtuous method of belief formation in the bad case where the calculator generates a false answer. The methods that she uses are equally globally reliable with respect to the same field of propositions and with respect to the same set of circumstances. Paige's belief is based on perception and inference in both cases. Finally, the circumstances where Paige forms her belief in the bad case are within the set of circumstances with respect to which the virtuous method that she uses in the

[12] Greco (2016) advocates a version of the safety condition that is relativized to the subject's cognitive virtues. He motivates the safety condition by arguing that the satisfaction of a properly formulated virtue-theoretic condition entails the satisfaction of a properly formulated safety condition.

[13] Epistemologists, such as Carter (2016), Gaultier (2014), Greco (2007), Littlejohn (2014), Turri (2016) and Pritchard (2012), who think that knowledge requires the satisfaction of both a virtue-theoretic condition and a safety condition, should be inclined to accept SAFETY^V as a necessary condition for knowledge. I (forthcoming) have argued elsewhere that virtue epistemologists should reject the idea that in cases of knowledge the cognitive success that the subject attains is creditable to the subject's cognitive abilities and embrace SAFETY^V instead.

actual world is globally reliable. After all, the lighting conditions are the same, she punches in the calculation correctly in both cases, she is sober and equally attentive in both cases, and so on. Therefore, she uses the same virtuous method of belief formation in both the good and the bad case. Given that there are nearby possible worlds where she gains a false belief by using the same virtuous method of belief formation that she uses in the actual world, she does not satisfy SAFETY^V.

SAFETY^V delivers the same verdict in WATER, BOWL OF FRUIT, and in THERMOMETERS. In all of these cases, the subject's belief is virtuously formed both in the good and the bad case, but given that the subject could easily have been in the bad case, where her belief is false, she does not satisfy SAFETY^V. Before moving onwards, let me show how SAFETY^V handles a classic case offered by Goldman.

DACHSHUND: Oscar sees Dack the dachshund before him and forms (non-inferentially) the belief that a dog is before him. However, if the dachshund had not been there, a wolf would have been there instead, which Oscar would have misclassified as a dog, since Oscar has a tendency to mistake wolves for dogs. (Goldman 1976, p. 779)

It seems obvious that Oscar knows that a dog is before him when he sees a dachshund. However, if Oscar uses the same virtuous method of belief formation in the actual and counterfactual case, then his belief is unsafe. Let us test whether Oscar uses the same virtuous method of belief formation in both cases.

Condition (ii) is satisfied since Oscar gains his belief by exercising his faculty of vision in both cases. What about condition (i)? Given that Oscar knows that a dog is before, him in the good case, his belief is virtuously formed (he is globally reliable in telling whether something that looks like a dachshund is a dog in the circumstances in which he forms his belief). If we understand the field of propositions along the lines suggested by Sosa, the propositions that Oscar believes in the good and bad cases belong to the same field of propositions, since they are both of the form [that is a dog]. Given that nothing in the case description of DACHSHUND prevents us from assuming that the method of belief-formation that Oscar uses in the bad case is equally reliable with respect to the same field of propositions and same range of circumstances as the method he uses in the good case, condition (i) can be satisfied by Oscar. However, condition (iii) is not satisfied by Oscar. The relevant set of circumstances is not the same in the actual and counterfactual scenario. In the actual world, Oscar forms his belief on the basis of seeing a sausage-like, furry, floppy-eared, four-legged creature that has features characteristic of dachshunds. In the counterfactual scenario, Oscar forms his belief on the basis of entirely different visual cues. The creature before him is considerably larger, has pointy ears, its body is anything but sausage-shaped, etc. Given that Oscar has a tendency to misclassify wolves as dogs, the set of circumstances in which he is in the counterfactual scenario do not belong to the set of circumstances regarding which the method that he uses in the actual world is globally reliable. Therefore, condition (iii) is not satisfied. Since condition (iii) is not satisfied, there are no nearby possible worlds where Oscar believes that a dog is before him by using the same virtuous method of belief formation that he uses in the actual world, when his belief is false. Therefore, Oscar's belief is virtuously safe and SAFETY^V delivers the correct verdict in DACHSHUND.

But does SAFETY^V restrict the relevant set of propositions in an adequate way? How, for example, should the field of propositions be defined in the case of good eyesight? Presumably, if one has good eyesight, then one is able to distinguish medium-sized objects from other medium-sized objects, to locate medium-sized objects in one's visual field, to follow the movement of medium sized-objects, etc. The field of propositions will in that case consist of propositions like [the medium-sized object x is there], [that medium-sized object is an F].

Consider the barn façade case. Intuitively, Barney does not know that he is looking at a barn when he sees one in the barn façade county. SAFETY^V is able to accommodate this intuition. But suppose that Barney, while in the barn façade-county, happens to look at a rock, and forms the true and justified belief that [that's a rock]. Does Barney know that it is a rock? Intuitively, the answer is 'yes'. However, the rock is a medium-sized object, and Barney could easily have formed a false belief ([that's a barn]) about another medium-sized object. There is no reason to suppose that Barney's beliefs, [that's a rock] and [that's a barn], are not formed in an equally globally reliable way. Furthermore, it seems that the circumstances, in which Barney forms his belief about the rock, are in the set of circumstances with respect to which his eyesight is globally reliable in the case where he forms a belief about the barn. Therefore, Barney could easily have formed a false belief via the same virtuous method of belief formation that he uses in the actual situation. Barney does not know that [that's a rock] according to SAFETY^V. The same line of reasoning can be applied to the BOWL OF FRUIT case. The fact that the bowl contains convincing fruit replicas does not undermine Sofia's knowledge that a bowl is on the table. SAFEY^V is unable to accommodate this fact, because Sofia uses the same virtuous method in the case where she believes falsely that [that's a persimmon] and in the case where she believes truly that [the bowl is on the table].

Restricting the relevant set of propositions solely in terms of the method of belief formation that the subject uses in the actual world is not good enough. We need to restrict the relevant set of propositions through other means as well, as Williamson did. However, Williamson's way of restricting the relevant set of propositions was not satisfactory. Let me offer a more promising way of restricting the relevant set of propositions.

In order to accomplish this we need to consider what the subject has to be safe from in order to have knowledge. I suggest that the relevant set of propositions should be restricted in terms of the agent's subject matter of inquiry. In order to be properly safe from error with respect to p, it has to be the case that one could not easily have ended up with a false belief in one's inquiry whether p. Alternatively, if S knows that p, there is a question Q to which p is a correct answer and S could not easily have believed in a false answer to Q.[14] Restricting the relevant set of propositions in this way yields correct verdicts about the cases that we have considered.

[14] Ball has recently proposed that "a subject S's belief that p is safe if, and only if, S could not easily have believed a false answer to the question Q to which p is saliently an answer" (2016, p. 60). He recognizes that the safety condition should be relativized to a question to which the proposition believed in the actual world is an answer. However, he does not commit himself to any view as to how the relevant question is to be determined nor does he define the safety condition in detail. While his view lacks some details that the view offered here possesses I take it that his view is compatible with the account I propose and able to deal with cases featuring necessary truths.

In BOWL OF FRUIT Sofia is interested in knowing what is in the bowl. The relevant set of propositions will be comprised of propositions that are about the contents of the bowl, such as [that's an apple], [that's an orange], and [that's a persimmon]. If Sofia had believed in the actual world that [the bowl is on the table], the subject matter of her inquiry would have been a different one than the one she had in the original case. The fact that she could easily have formed a false belief about the contents of the bowl does not mean that she does not know that there is a bowl in front of her.[15] In MALFUNCTION Paige could easily have ended up with a false belief in her inquiry of what is the product of 12×13. Given that the calculator is generating answers at random there are plenty of nearby possible worlds where Paige ends up with a false belief about the product of 12×13 or some other equation.[16]

But how do we determine the agent's subject matter of inquiry in a given case? Rough guidelines are the best that I can offer on this score. The practical interests of the subject help to determine the subject matter of inquiry. In evaluating what is the subject matter of inquiry in a given case we should ask ourselves "What is the agent trying to find out?" This is often implicit in a case description and the belief that the subject formed in the actual world will reveal the subject matter of inquiry. Of course, sometimes the agent is not trying to find out anything at all. If I hear someone say in the corridor that "It is raining" and passively accept this piece of testimony I might be able to know that [It is raining] if everything goes well. In such a case it seems that I did not try find out anything. But this does not rule out the possibility that a subject matter of inquiry could be attributed to me. The context of knowledge attribution can aid us to determine the agent's subject matter of inquiry. The practical interests of the subject or the epistemic community help us to see what kind of beliefs the subject could easily have gained in her epistemic endeavours. Sometimes the subject matter will be harder to discern, but that is to be expected. Even a good definition of the safety condition might not be very easy to use in every case.

The fact that the agent's subject matter of inquiry is partially determined through the practical interests and the context of knowledge attribution does not make the proposed safety condition relativistic. True, a subject who has not managed to narrow down her

[15] But could we not claim that Sofia's subject matter of inquiry is "are there any apples in the bowl?" If that was her subject matter of inquiry she could not easily have ended up with a false belief in her inquiry and hence would satisfy the proposed safety condition (I would like to thank an anonymous referee for raising this objection). In order to make it salient that Sofia's subject matter of inquiry is restricted to this very particular question we would have to add a lot of details to the case. It could not, for instance, be the case that Sofia is about make a fruit salad, since in that case she would be interested in finding fruits. Crucially, if we fill in the missing details of the case so that Sofia's subject matter of inquiry cannot be thought to be "what is in the bowl", but is saliently "are there any apples in the bowl" there is pressure to think that Sofia gains knowledge in the case after all. Greco (2012, p. 23), for example, has argued that cases of environmental luck, such as the barn façade case, do not seem to be cases of ignorance if the practical task that the subject sets out to complete is restricted in such a way that the subject is guaranteed to form a true belief in completing the task. Greco emphasizes that most knowledge claims are not so narrowly associated with a specific task, and hence in most contexts of knowledge attribution the subject in the barn façade county does not have knowledge.

[16] The difference between the field of propositions F and subject matter of inquiry Q is that F is comprised of propositions that the subject could form a belief in by exercising the epistemic virtue that she uses in the actual world to form her belief and Q is comprised of the propositions that are possible answers to the question that the subject aims to answer.

subject matter of inquiry is required to do more than a subject who has managed to narrow down her subject matter of inquiry, in order to have knowledge, but this is to be expected. As Hookway (1996, p. 14) points out, asking the right questions is as important as answering them. By asking the right kind of questions one is better situated to gain knowledge.

Given that the proposed version of the safety condition is defined by reference to the agent's subject matter of inquiry or the question that the subject is trying to solve, one might think that the view entails that knowledge is a contrastive notion, as has been argued by Schaffer (2004, 2007). According to contrastivists, "S knows that p" means that S knows that p rather than Q, where Q is a set of contrast propositions consisting of possible answers to the question that is being discussed (Schaffer 2007, pp. 400–401).[17] Depending on what values are assigned to Q, a subject might know or fail to know in a given situation. For example, in Dretske's painted mule case, the subject knows that there is a zebra in the pen if Q consists of propositions such as [there is a giraffe in the pen], [there is a rhino in the pen], but not if Q contains propositions such as [there is a cleverly painted mule in the pen], because the subject cannot discriminate cleverly painted mules from zebras.[18] The safety condition advocated here does not entail that knowledge is contrastive; it entails that if a subject knows that p, there is a question Q to which p is a correct answer and S could not easily have believed in a false answer to Q. If the relevant question is "What animal is in the pen?" one possible answer will be that [there is a cleverly painted mule]. The fact that Q contains this proposition does not mean that the subject does not know that there is a zebra in the pen, as long as there are no nearby possible worlds where there is a cleverly painted mule in the pen.

Finally, it is worth emphasizing that restricting the relevant set of propositions in terms of subject matters of inquiry is not *ad hoc*. After all, it is extremely plausible that in order to know that p one must be safe from error in one's inquiry whether p. Restricting the set in terms of subject matters of inquiry is also quite informative, at least compared to Williamson's way of trying to restrict the relevant set of propositions.

Taking these remarks into account we get the following globalized version of SAFETY^V:

GLOBAL SAFETY: S's belief that p, which belongs to her subject matter of inquiry Q, is safe if and only if:

(i) in most nearby possible worlds, and in all of the very closest possible worlds, where S believes in a proposition that belongs to Q via the same virtuous method V that S uses in the actual world, S's belief is true.

6 Concluding remarks

GLOBAL SAFETY is able to deal with cases featuring necessary truths without succumbing to scepticism. Moreover, GLOBAL SAFETY is preferable to SAFETY^P and

[17] See also Karjalainen and Morton (2003).

[18] See Dretske (1970, p. 1016) for the case.

SAFETY^W in several respects. Firstly, the way in which the method of belief formation is individuated in GLOBAL SAFETY does not deliver wrong verdicts in cases like BOWL OF FRUIT, THERMOMETERS and WATER. Secondly, restricting the methods of belief formation to virtuous methods of belief formation allows the safety theorist to offer a promising solution to the generality problem. Finally, restricting the relevant set of propositions in terms of subject matters of inquiry is more informative than restricting it in terms of propositional closeness. Once the safety condition is properly understood, cases featuring necessary truths pose no problems for the idea that in order to know, one has to be safe from error.

Acknowledgements I would like to thank audiences at Uppsala, Tampere and Helsinki where this paper was presented. Thanks also to Matti Eklund, Jani Hakkarainen, Sebastiano Moruzzi, Jaakko Kuorikoski, Markus Lammenranta, Tuukka Tanninen, Peter Schulte and two anonymous referees at *Synthese* whose comments helped me improve this paper.

References

Ball, B. (2016). Knowledge, safety, and questions. *Unisinos Journal of Philosophy*, *17*(1), 58–62.
Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research*, *88*(2), 289–313. doi:10.1111/j.1933-1592.2011.00564.x.
Bogardus, T., & Marxen, C. (2014). Yes, safety is in danger. *Philosophia*, *42*, 321–334.
Broncano-Berrocal, F. (2014). Is safety in danger? *Philosophia*, *42*(1), 63–81. doi:10.1007/s11406-013-9467-9.
Carter, J. A. (2016). Robust virtue epistemology as anti-luck epistemology: A new solution. *Pacific Philosophical Quarterly, 97*(1), 140–155. doi:10.1111/papq.12040.
Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, *58*(1), 7–19.
Coffman, E. J. (2007). Thinking about luck. *Synthese*, *158*(3), 385–398. doi:10.1007/s11229-006-9046-8.
Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, *89*, 1–29.
Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy*, *67*(24), 1007–1023. doi:10.2307/2024710.
Gaultier, B. (2014). Achievements safety and environmental epistemic luck. *Dialectica*, *68*(4), 477–497. doi:10.1111/1746-8361.12081.
Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, *73*, 771–791.
Goldman, A. (1986). *Epistemology and cognition*. Cambridge: Harvard University Press.
Greco, J. (2007). Worries about Pritchard's safety. *Synthese*, *158*(3), 299–302. doi:10.1007/s11229-006-9040-1.
Greco, J. (2010). *Achieving knowledge: A virtue theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
Greco, J. (2012). A (Different) virtue epistemology. *Philosophy and Phenomenological Research*, *85*(1), 1–26. doi:10.1111/j.1933-1592.2011.00567.x.
Greco, J. (2016). Knowledge, virtue, and safety. In M. A. Fernández Vargas (Ed.), *Performance epistemology—Foundations and applications* (pp. 51–61). Oxford: Oxford University Press.
Hawthorne, J. (2006). *Metaphysical Essays*. Oxford: Clarendon Press.
Hirvelä, J. (forthcoming). On virtue, credit and safety. *Grazer Philosophische Studien*.
Hookway, C. (1996). The presidential address: Questions of context. *Proceedings of the Aristotelian Society*, *96*, 1–16.
Karjalainen, A., & Morton, A. (2003). Contrastive knowledge. *Philosophical Explorations*, *6*(2), 74–89. doi:10.1080/10002003058538741.
Littlejohn, C. (2014). Fake barns and false dilemmas. *Episteme—A Journal of Individual and Social Epistemology*, *11*(4), 369–389. doi:10.1017/epi.2014.24.
Luper, S. (2006). Restorative rigging and the safe indication account. *Synthese*, *153*(1), 161–170. doi:10.1007/s11229-005-6399-3.
Neta, R., & Rohrbaugh, G. (2004). Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, *85*(4), 396–406. doi:10.1111/j.1468-0114.2004.00207.x.

Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, *158*(3), 277–297. doi:10.1007/s11229-006-9039-7.
Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, *109*(3), 247–279.
Pritchard, D. (2015). The modal account of luck. In D. Pritchard & L. Whittington (Eds.), *The philosophy of luck*. London: Wiley-Blackwell.
Pritchard, D. (2016). Anti-luck virtue epistemology and epistemic defeat. *Synthese*, pp. 1–13, doi:10.1007/s11229-016-1074-4.
Riggs, W. (2007). Why epistemologists are so down on their luck. *Synthese*, *158*(3), 329–344. doi:10.1007/s11229-006-9043-y.
Roland, J., & Cogburn, J. (2011). Anti-luck epistemologies and necessary truths. *Philosophia*, *39*(3), 547–561. doi:10.1007/s11406-010-9295-0.
Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, *119*(1), 73–103. doi:10.1023/B:PHIL.0000029351.56460.8c.
Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, *75*(2), 383–403. doi:10.1111/j.1933-1592.2007.00081.x.
Sosa, E. (1991). *Knowledge in perspective—selected essays in epistemology*. Cambridge: Cambridge University Press.
Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, *13*, 141–154.
Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford: Oxford University Press.
Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives*, *24*(1), 465–475. doi:10.1111/j.1520-8583.2010.00200.x.
Turri, J. (2016). Knowledge as achievement, more or less. In M. A. Fernández Vargas (Ed.), *Performance epistemology—foundations and applications* (pp. 124–134). Oxford: Oxford University Press.
Williamson, T. (1994). *Vagueness*. London: Routledge.
Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
Williamson, T. (2009a). Reply to Alvin Goldman. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 305–312). Oxford: Oxford University Press.
Williamson, T. (2009b). Reply to John Hawthorne and Maria Lasonen-Aarnio. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 313–329). Oxford: Oxford University Press.
Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.