

Knowledge, Safety, and Questions

Brian Ball¹

ABSTRACT

Safety-based theories of knowledge face a difficulty surrounding necessary truths: no subject could have easily falsely believed such a proposition. Failing to predict that ill-grounded beliefs in such propositions do not constitute knowledge, standard safety theories are therefore less informative than desired. Some have suggested that the subjects at issue could easily have believed some related false proposition; but they have given no indication as to what makes a proposition related. I suggest a solution to this problem: a belief is safe iff its subject could not easily have believed a false answer to the same question.

Keywords: epistemology, knowledge, problem of necessary truths, questions, safety.

Knowledge and Safety

A number of authors have advocated a safety constraint on knowledge. Thus, Ernest Sosa, for example, claims that “[i]n order to [...] constitute knowledge a belief must be safe” (1999, p. 142); or, in other words, “[n]o belief constitutes knowledge unless safe” (1999, p. 143). Similarly, Timothy Williamson says that “[i]f one knows [...] one’s belief is safely true” (2000, p. 147). Other examples abound in the literature (e.g. Pritchard, 2007; Sainsbury, 1997).

Clearly, these theorists think that the safety of a belief is *necessary* for knowledge. But what is safety? More specifically, what is it for a belief to be safe, or safely true? The notion is a modal one. Sosa says that a subject *S*’s belief that *p* is safe if, and only if, “*S* would not believe that *p* without it being the case that *p*” (1999, p. 142). Williamson suggests that safety consists in “the avoidance of false belief at close worlds” (2000, p. 149), and that one’s belief is safely true if, and only if, “one could not easily have been wrong in a similar case” (2000, p. 147). And Jonathan Jenkins Ichikawa and Matthias Steup propose that to say that *S*’s belief that *p* is safe is to say that “[i]n all nearby worlds where *S* believes that *p*, *p* is not false” (2014, section 5.2), but is, presumably, true instead.² For the sake of determinacy, let us stick with this last account, which is, I believe, quite standard.

Safety theorists have not, in general, claimed that this condition (possessing a safe belief that *p*), which entails both truth (because the actual world is nearby to itself) and belief (since one cannot have a safe belief without having a belief), is *sufficient* for knowledge (that *p*).³ So if it should turn out that there are beliefs that are safe in this sense but which do not, intuitively, constitute knowledge, that is not, in and of itself, a problem for the view. And indeed, it seems relatively easy to come up with examples in which beliefs are safe in this sense yet don’t constitute knowledge. Suppose, for instance, that Crazy Jailbird is in a maximum security prison, such that he couldn’t easily fail to be in prison tomorrow night. However, Crazy believes utterly erroneously that he has friends on the outside who are conspiring to free him; moreover, he believes that a certain coin, which is in fact

¹ Balliol College, Oxford, OX1 3BJ, United Kingdom. New College of the Humanities, 19 Bedford Square, London, WC1B 3HH, United Kingdom. E-mail: brian.ball@philosophy.ox.ac.uk; brian.ball@nchlondon.ac.uk.

² Pritchard (2007, p. 71) speaks instead of the belief in question being true in *most* nearby worlds in which it is held. This weakening does nothing to alleviate the problem of necessary truths discussed below.

³ For instance, Sainsbury says that “if you know, you could not easily have been wrong. No equivalence is claimed: the absence of easy possibility [of error] is only claimed to be a necessary condition for knowledge” (1997, p. 908).

fair, serves as an oracle. He comes to the view that if the next toss of his coin lands heads he will be freed tomorrow, but that he will not be freed until later in the week if the coin lands tails. He tosses the coin, it lands tails, and he comes to the view that he will be in prison tomorrow night. If this is the only reason Crazy believes he will be in prison tomorrow night, then he does not know that he will be, even though his belief is safe.

So the idea is not that knowledge consists in safe belief: rather, it seems that safety is a consequence of knowledge. More specifically, the safety of a knowledgeable belief might be thought to follow from the fact that knowledge must have an appropriate doxastic-cum-epistemic basis or ground. If a belief has such a basis—which, if it is to be knowledge it must—then, the argument goes, if its object were false, that basis for belief would not have been available, and the belief would not have been held. So the belief in question could not easily have been falsely held, and knowledgeable beliefs are safe.

That this is the central idea underlying safety theories can serve to explain the following otherwise potentially puzzling claim of Williamson's:

In many cases, someone with no idea of what knowledge is would be unable to determine whether safety obtained. [...] One may have to decide whether safety obtains [in such cases] by first deciding whether knowledge obtains, rather than vice versa (2009, p. 305).

If safety theorists such as Williamson were proposing to reduce knowledge to safely true belief, this might be thought to be problematic; but if instead safety is held to be a consequence of knowledge, then it can simply be maintained that certain possibilities are more similar to the actual case (and therefore could easily have obtained) than others (which could not) because they hold fixed whether knowledge obtains in those cases, and so they hold fixed the actual grounds for belief.

Nonetheless, even if safety is not thought to be sufficient for knowledge, the claim that safety is necessary for knowledge yields determinate predictions, and so is informative. For it is possible to argue, and in some sense explain the fact, that a belief does not constitute knowledge by showing that it is not safe. If Jailbird could easily have been paroled tomorrow morning, his belief that he will be in prison tomorrow night would not constitute knowledge since it could easily have been false, not being connected to its truth in an appropriate manner.

The Problem of Necessary Truths

There is, however, a problem. Any belief in a necessary truth is safe. While this fact does not present the theory with

any counterexamples, it does make the safety condition, as formulated above, utterly trivial and uninformative in these cases. We cannot argue that a belief in a necessary truth is not knowledge on the grounds that it isn't safe. For instance, if Crazy Mathematician thinks her actually fair coin is an oracle, then she might come to the view that Fermat's Last Theorem is true if, and only if, the next toss of her coin lands heads: if it then does land heads and she comes to believe that Fermat's Last Theorem holds, her belief does not constitute knowledge; yet we cannot argue that it does not by appeal to the safety condition articulated above, for her belief *couldn't* have been false, and so, *a fortiori*, could not easily have been false. This problem of necessary truths might be thought to restrict the interest of the safety requirement on knowledge.⁴

Williamson (2000, p. 181-182) responds to this problem by suggesting that what matters for the safety of a belief is not whether that very belief could easily have been false but whether the subject could easily have had a false belief in a related proposition.⁵ This, it seems to me, is correct in so far as it goes. But it is vague. Which propositions count as related?

My proposal, roughly speaking, is that a proposition is related to the proposition that *p* if, and only if, it is an answer to the same question. One way to fill this out is to say that a belief that *p* is safe if, and only if, the subject could not easily have had a false belief on the question whether *p*. As we shall see, if this specific proposal is adopted, problematic cases will remain: but perhaps a slightly more general idea will suffice; and in any case, even the specific proposal yields a non-trivial constraint on our knowledge of necessities. But we will come to all of this in due course.

Questions

First, we must ask: What is a question? It will be useful to approach this issue via the distinct but related question, What is a proposition? Propositions are quite familiar to philosophers. They are the objects of certain speech acts, e.g. the declarative speech act of assertion; they are what we assert when we assert something. They also occur semantically embedded in indirect discourse constructions such as "Amy said that it is snowing," and in certain attitude ascriptions, e.g. "John believes that it is raining." Finally, they are the objects of certain psychological attitudes, such as belief—namely, the *propositional attitudes*.

More substantively, there has been some dispute about what propositions are. For instance, there is an internal disagreement amongst those who think of them as structured entities about what the nature of their constituents might be, with neo-Russellians holding that they comprise objects, properties, and relations as parts, and neo-Fregeans maintain-

⁴ Indeed, Pritchard (2007, p. 71) explicitly, though unnecessarily, restricts attention to contingent propositions in his formulation of the safety requirement on knowledge.

⁵ See also Sainsbury (1997, p. 908-909).

ing that the constituents in question are rather modes of presentation of such worldly entities: however, those working in formal semantics have found it useful to think of propositions as (unstructured) sets of possible worlds; and in any case, everyone can at least agree that they *determine* sets of worlds.

Similarly, then, a question is the object of a certain kind of speech act, e.g. the interrogative act of asking: it is what we ask when we ask a question. It is also, semantically speaking, the object of the verb in certain linguistic constructions, e.g. "John wondered who took the last cookie," and "Mary asked whether it was raining." And Jane Friedman (2013) has recently argued that questions are the objects of certain kinds of psychological attitudes—attitudes such as wondering, which she calls *interrogative attitudes*.

Formal theorists (e.g. Hamblin, 1973) have found it useful to think of questions as sets of their (complete) answers, that is, as sets of propositions that (completely) answer them. Thus, if propositions are sets of possible worlds, questions are families (or sets) of (certain) such sets; and in any case, everyone can agree that they *determine* such sets of sets of worlds. Thus, we may say that questions are, or determine, partitions of the domain of possible worlds.⁶

Some answers are complete and others are incomplete. Suppose I host a party, that Amy, Bob, and Charlie come, and that Doug and Emily do not. (And let's suppose that there are no other people.) Then the question, Who came to the party? can be answered completely by saying that Amy, Bob, and Charlie did. Another possible complete answer is that Charlie and Doug did. This alternative answer is not a correct answer, but it is a possibly correct, complete answer. By contrast, the proposition that Amy came to the party correctly answers the question, but does not answer it completely.

A Solution

With this background in place, we can now say what it is for a belief to be safe. A first, simple proposal is that *S*'s belief that *p* is safe if, and only if, *S* could not easily have had a false belief on the question of whether *p*—that is, the question whose two complete answers are *p* and *not p*. This deals with some instances of the problem of necessary truths, for instance, the case of the Crazy Mathematician discussed above. The subject's belief in that case is safe in the Ichikawa-Steup sense above, but intuitively it does not constitute knowledge. Yet the new account of safety readily accounts for the failure of knowledge in this case, since the Crazy Mathematician could easily have had a false belief on the question of whether Fermat's Last Theorem is true; indeed, she would have done had the coin she tossed landed tails (which it could easily have

done). Thus, her belief is not safe in the new sense; and since its being so is necessary for it to constitute knowledge, she does not know Fermat's Last Theorem.

However, there remain problematic cases. Consider the following example, due to Roland and Cogburn, which involves Sam, whose calculator is broken

*so that it always tells the user that whatever number entered is prime. Sam uses his calculator to randomly check whether or not [some prime number *p*] is prime, and it answers affirmatively. As a result, Sam forms the true belief that [*p*] is a prime number (2011, p. 550).*

In this case of the Broken Calculator, Sam could not easily have formed a false belief on the question of whether *p* is prime, since he could not easily have formed the belief that it is not prime: accordingly, his belief is safe, not only in the Ichikawa/Steup sense, but also in the sense just articulated; and Broken Calculator therefore shows that the safety constraint so construed is not sufficient for knowledge.

There are two things that might be said in response to this case. The first, more ambitious reply, begins by noting that the proposition that *p* is prime does not only answer the question whether *p* is prime; it is also a partial answer to the question, Which numbers are prime? Moreover, Sam could easily have had a false belief on this question: after all, he selects *p* as the number to check at random; and had he decided to check some composite number *c* instead, he would have formed a false belief in (partial) answer to the question of which numbers are prime—namely, the false belief that *c* is prime.

This suggests that a more subtle safety condition on knowledge might be articulated as follows: 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. Clearly, though, this suggestion does not determine precisely which question this is; so it does not yet resolve the vagueness of the Williamsonian proposal mentioned above. Nonetheless, there are a number of ways in which the proposal might be made more concrete, thereby resolving at least some of the underlying vagueness; and the appeal to a salient question might provide a means of unifying the various propositions that are relevant to the safety of the belief at issue in a given case.

The second, more modest reply is to note that our aim was to find an informative constraint on knowledge of necessary truths; and, as we have seen in connection with Crazy Mathematician, the simple proposal that a subject *S*'s belief

⁶ Some questions have presuppositions—for instance, the question of whether you have stopped beating your wife. These may not appear to partition the whole space of possible worlds: some (ugly) worlds are ones in which you continue to beat your wife, and some (bad) worlds are ones in which you have stopped; but there are also (good) worlds in which you have never beaten your wife (perhaps because you don't have one). Formally, however, we can treat this final case as just one more cell in the partition of worlds that constitutes the question.

that *p* is safe if and only if *S* could not easily have had a false belief on the question of whether *p* provides just that. Thus, in particular, if the case of Crazy Jailbird does not undermine safety-based approaches to knowledge, then the case of the Broken Calculator shouldn't be taken to undermine the proposal advanced here either. On the current proposal, safe belief is not sufficient for knowledge, but it is both necessary and non-trivial, even when it comes to our knowledge of necessities—and that is exactly what was sought.⁷

Concluding Remarks

I am not the first to propose that the theory of knowledge may benefit from the deployment of the notion of a question. Jonathan Schaffer (2007) has argued that knowledge is a three-place relation between a subject, a proposition, and a question.⁸ If he is right, then (assuming safety is a necessary condition on knowledge) that might explain why we need to appeal to questions when explaining the notion of safety. Nevertheless, my proposal is different from—indeed, less committal than—Schaffer's, in two specific (though not entirely unrelated) respects. First, even on the ambitious proposal, I am not suggesting anything about the semantics in general, or the logical form in particular, of knowledge attributions. Accordingly, while the question that is salient might be determined by some feature of the context of speech, this is by no means required; it might, for instance, be fixed by some feature of the situation in which the subject finds him or herself.⁹ Second, my proposal, whether modest or ambitious, is concerned directly with safety, not knowledge; it therefore has no immediate implications for those who reject safety-based approaches to epistemology. Accordingly, those who reject Schaffer's position may nonetheless find something of value in what I have said here.

Other theorists have appealed tacitly to questions, while failing to do so explicitly. Thus, Nozick, for instance, in characterizing his sensitivity-based account of knowledge, suggests that a subject who knows that *p* is such that she would believe that *p* if she were to “have a belief whether (or not) *p*” (1981, p. 179) in a case in which *p*; and she would not have the belief that *p* in the case in which not *p*. But the quoted

phrase is not quite grammatical as it stands; what's needed for full acceptability is the insertion of (e.g.) “on the question (of)” between “belief” and “whether.” With questions now being granted the philosophical respectability long afforded to propositions, we can have no qualms about explicitly recognizing their role in the theory of knowledge.¹⁰

In any case, the notion of a question has, until recently, been relatively neglected in philosophical theorizing. I have found a new application for it—in explicating the notion of safety in such a way that safety theorists of knowledge have a (relatively) precise and non-trivial theory of our knowledge of necessary truths. This not only constitutes an advance on this specific topic; it also provides some further, if modest, inductive evidence of the theoretical utility of questions more generally.¹¹

References

- DEROSE, K. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, **52**(4):913-929. <http://dx.doi.org/10.2307/2107917>
- FRIEDMAN, J. 2013. Question-Directed Attitudes. *Philosophical Perspectives*, **27**(1):145-174. <http://dx.doi.org/10.1111/phpe.12026>
- HAMBLIN, C.L. 1973. Questions in Montague English. *Foundations of Language*, **10**(1):41-53.
- HAWTHORNE, J. 2004. *Knowledge and Lotteries*. Oxford, Oxford University Press, 216 p.
- ICHIKAWA, J.J.; STEUP, M. 2014. The Analysis of Knowledge. In: E.N. ZALTA (ed.), *The Stanford Encyclopedia of Philosophy* Spring 2014 Edition. Available at: <http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>. Accessed on: 05/09/2016.
- LEWIS, D. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4):549-567. <http://dx.doi.org/10.1080/00048409612347521>
- NOZICK, R. 1981. *Philosophical Explanations*. Cambridge, Harvard University Press, 784 p.
- PRITCHARD, D. 2007. *Epistemic Luck*. Oxford, Oxford University Press, 304 p.
- ROLAND, J.; COGBURN, J. 2011. Anti-Luck Epistemologies and Necessary Truths. *Philosophia*, **39**(3):547-561. <http://dx.doi.org/10.1007/s11406-010-9295-0>

⁷ Roland and Cogburn construe safety theorists as committed to the presence of a safely true belief as being both necessary and sufficient for knowledge, and therefore regard Broken Calculator as a counterexample. As we have seen, Sainsbury, Sosa, and Williamson, for instance, incur no such commitment.

⁸ This builds on his earlier (Schaffer, 2004) contrastivist view that it is a three-place relation between a subject, a proposition *p*, and a second, contrast proposition *q*. This might be taken to be roughly equivalent to the view that to know that *p* is to know *p* as the answer to the question of whether *p* or *q*, and thereby assimilated to the later (Schaffer, 2007) view.

⁹ Cf. the difference between contextualist (DeRose, 1992; Lewis, 1996) and subject sensitive invariantist (Hawthorne, 2004; Stanley, 2005) accounts of knowledge.

¹⁰ In fact, Nozick (1981, p. 186-187) also uses this same tacit appeal to questions to address the problem of necessary truths, as it confronts his sensitivity account of knowledge. I claim no great originality for my solution to the problem, but only for the explicit recognition of the form that that solution takes.

¹¹ Thanks to audiences at the Joint Session 2015 and the Philosophy Society at New College of the Humanities for feedback on earlier versions.

- SAINSBURY, R.M. 1997. Easy Possibilities. *Philosophy and Phenomenological Research*, **57**(4):907-919.
<http://dx.doi.org/10.2307/2953809>
- SCHAFFER, J. 2004. From Contextualism to Contrastivism. *Philosophical Studies*, **119**(1):73-103.
<http://dx.doi.org/10.1023/B:PHIL.0000029351.56460.8c>
- SCHAFFER, J. 2007. Knowing the Answer. *Philosophy and Phenomenological Research*, **75**(2):383-403.
<http://dx.doi.org/10.1111/j.1933-1592.2007.00081.x>
- SOSA, E. 1999. How to Defeat Opposition to Moore. *Philosophical Perspectives*, **33**(s13):141-153.
<http://dx.doi.org/10.1111/0029-4624.33.s13.7>
- STANLEY, J. 2005. *Knowledge and Practical Interests*. Oxford, Oxford University Press, 191 p.
<http://dx.doi.org/10.1093/0199288038.001.0001>
- WILLIAMSON, T. 2000. *Knowledge and its Limits*. Oxford, Oxford University Press, 352 p.
- WILLIAMSON, T. 2009. Reply to Alvin Goldman. In: P. GREENOUGH; D. PRITCHARD (eds.), *Williamson on Knowledge*. Oxford, Oxford University Press, p. 307-312.

Submitted on February 9, 2016

Accepted on June 17, 2016